arXiv:1401.5492v4 [stat.ML] 8 Dec 2015

Alternating direction method of multipliers for penalized zero-variance discriminant analysis

Brendan P.W. Ames [*] Mingyi Hong [†]

December 9, 2015

Abstract

We consider the task of classification in the high dimensional setting where the number of features of the given data is significantly greater than the number of observations. To accomplish this task, we propose a heuristic, called sparse zero-variance discriminant analysis (SZVD), for simultaneously performing linear discriminant analysis and feature selection on high dimensional data. This method combines classical zero-variance discriminant analysis, where discriminant vectors are identified in the null space of the sample within-class covariance matrix, with penalization applied to induce sparse structures in the resulting vectors. To approximately solve the resulting nonconvex problem, we develop a simple algorithm based on the alternating direction method of multipliers. Further, we show that this algorithm is applicable to a larger class of penalized generalized eigenvalue problems, including a particular relaxation of the sparse principal component analysis problem. Finally, we establish theoretical guarantees for convergence of our algorithm to stationary points of the original nonconvex problem, and empirically demonstrate the effectiveness of our heuristic for classifying simulated data and data drawn from applications in time-series classification.

1 Introduction

A standard technique for supervised classification is to perform dimensionality reduction to project the data to a lower dimensional space where the classes are well separated, and

[*]Department of Mathematics, The University of Alabama, Box 870350, Tuscaloosa, AL 35487-0350, bpames@ua.edu

[†]Department of Industrial and Manufacturing Systems Engineering, Iowa State University, Black Engineering Building, Ames, IA, 50011, mingyi@iastate.edu

then classify the data in this lower dimensional space. For example, a classical approach of Fisher [21], [44], [26, Chapter 4], called *linear discriminant analysis (LDA)*, is to define a linear mapping to a lower dimensional space where the ratio of the sample between-class scatter and within-class scatter of the projected training data is maximized, and then assign each test observation to the class of the nearest projected centroid. When the training data is randomly sampled from one of two Gaussian distributions with common covariance matrix, this approach reduces to Bayes' rule.

Before proceeding further, we first define some notation and assumptions. We represent our given data as the matrix $\boldsymbol{X} \in \mathbf{R}^{n \times p}$; here, each row $\boldsymbol{x}_i \in \mathbf{R}^p$ of $\boldsymbol{X}$ represents a single observation consisting of p features and the data set contains n such observations. We assume that the data has been centered and normalized so that each feature has mean equal to 0 and variance equal to 1. Further, we assume each observation is labeled as belonging to exactly one of k classes, denoted $C_1, C_2, \ldots, C_k$. Considered as separate data sets, the mean and covariance of each class C_i may be approximated by the sample class-mean $\hat{\boldsymbol{\mu}}_i$ and covariance matrix $\hat{\boldsymbol{\Sigma}}_i$ given by

$$\hat{\boldsymbol{\mu}}_i = \sum_{j \in C_i} \frac{\boldsymbol{x}_j}{|C_i|} \qquad \hat{\boldsymbol{\Sigma}}_{\boldsymbol{i}} = \frac{1}{n} \sum_{j \in C_i} (\boldsymbol{x}_j - \hat{\boldsymbol{\mu}}_i)(\boldsymbol{x}_j - \hat{\boldsymbol{\mu}}_i)^T.$$

We estimate variability within classes using the *sample within-class covariance matrix* $\boldsymbol{W}$, defined as the sum of the sample class-covariance matrices

$$\boldsymbol{W} = \sum_{i=1}^{k} \hat{\boldsymbol{\Sigma}}_i = \frac{1}{n} \sum_{i=1}^{k} \sum_{j \in C_i} (\boldsymbol{x}_j - \hat{\boldsymbol{\mu}}_i)(\boldsymbol{x}_j - \hat{\boldsymbol{\mu}}_i)^T,$$

and estimate variability between classes using the *sample between-class covariance matrix*,

$$\boldsymbol{B} = \frac{1}{n} \sum_{i=1}^{k} |C_i| \hat{\boldsymbol{\mu}}_i \hat{\boldsymbol{\mu}}_i^T;$$

note that $\boldsymbol{B}$ is defined as a rescaling of the sample covariance matrix of the data set obtained by replacing each observation with the sample mean of its class. As such, the column space of $\boldsymbol{B}$ is spanned by the k linearly dependent vectors $\hat{\boldsymbol{\mu}}_1, \hat{\boldsymbol{\mu}}_2, \ldots, \hat{\boldsymbol{\mu}}_k$ and, thus, has rank at most $k - 1$.

As mentioned earlier, we would like to identify a projection of the rows of X to a lower dimensional space where the projected class means are well separated, while observations within the same class are relatively close in the projected space. To do so, LDA seeks a set of nontrivial loading vectors $\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{k-1}$, obtained by repeatedly maximizing the criterion

$$J(\boldsymbol{w}) = \frac{\boldsymbol{w}^T \boldsymbol{B} \boldsymbol{w}}{\boldsymbol{w}^T \boldsymbol{W} \boldsymbol{w}}. \tag{1.1}$$

The fact that $\boldsymbol{B}$ has rank at most $k-1$ suggests that there exist at most $k-1$ orthogonal directions $\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{k-1}$ such that the quadratic form $\boldsymbol{w}_i^T \boldsymbol{B} \boldsymbol{w}_i$ has nonzero value.

To perform dimensionality reduction using LDA, we identify the desired loading vectors $\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{k-1}$ by sequentially solving the optimization problem

$$\boldsymbol{w}_i = \underset{\boldsymbol{w} \in \mathbf{R}^p}{\arg\max} \left\{ \frac{\boldsymbol{w}^T \boldsymbol{B} \boldsymbol{w}}{\boldsymbol{w}^T \boldsymbol{W} \boldsymbol{w}} : \boldsymbol{w}^T \boldsymbol{W} \boldsymbol{w}_j = 0 \; \forall \, j = 1, \ldots, i-1 \right\} \tag{1.2}$$

for all $i = 1, 2, 3, \ldots, k-1$. That is, $\boldsymbol{w}_i$ is the vector $\boldsymbol{W}$-conjugate to the span of $\{\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{i-1}\}$ that maximizes the LDA criterion (1.1). Noting that the criterion $J(\boldsymbol{w})$ is invariant to scaling, we may assume that $\boldsymbol{w}^T \boldsymbol{W} \boldsymbol{w} \leq 1$ and rewrite (1.2) as

$$\boldsymbol{w}_i = \underset{\boldsymbol{w} \in \mathbf{R}^p}{\arg\max} \left\{ \boldsymbol{w}^T \boldsymbol{B} \boldsymbol{w} : \boldsymbol{w}^T \boldsymbol{W} \boldsymbol{w} \leq 1, \; \boldsymbol{w}^T \boldsymbol{W} \boldsymbol{w}_j = 0 \; \forall \, j = 1, \ldots, i-1 \right\}. \tag{1.3}$$

Thus, finding the $k-1$ discriminant vectors $\{\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{i-1}\}$ is equivalent to solving the generalized eigenproblem (1.3). When the sample within-class covariance matrix $\boldsymbol{W}$ is nonsingular, we may solve (1.3) by performing the change of variables $\boldsymbol{z} = \boldsymbol{W}^{1/2} \boldsymbol{w}$. After this change of variables, we have $\boldsymbol{w}_i = \boldsymbol{W}^{-1/2} \boldsymbol{z}_i$, where

$$\boldsymbol{z}_i = \underset{\boldsymbol{z} \in \mathbf{R}^p}{\arg\max} \left\{ \boldsymbol{z}^T \boldsymbol{W}^{-1/2} \boldsymbol{B} \boldsymbol{W}^{-1/2} \boldsymbol{z} : \boldsymbol{z}^T \boldsymbol{z} \leq 1, \, \boldsymbol{z}^T \boldsymbol{z}_j = 0, \, j = 1, \ldots, i-1 \right\}.$$

That is, we may find the desired set of discriminant vectors by finding the set of nontrivial unit eigenvectors of $\boldsymbol{W}^{-1/2} \boldsymbol{B} \boldsymbol{W}^{-1/2}$ and multiplying each eigenvector by $\boldsymbol{W}^{1/2}$.

In the high dimension, low sample size setting, $p > n$, the sample within-class covariance matrix $\boldsymbol{W}$ is singular. Indeed, $\boldsymbol{W}$ has rank at most n, as it is a linear combination of n rank-one matrices. In this case, the change of variables $\boldsymbol{z} = \boldsymbol{W}^{1/2} \boldsymbol{w}$ used to compute the discriminant vectors is not well-defined. Moreover, the objectives of (1.2) and (1.3) can be made arbitrarily large if there exists some vector $\boldsymbol{w}$ in the null space of $\boldsymbol{W}$ not belonging to the null space of $\boldsymbol{B}$. To address this singularity issue, and to increase interpretability of the discriminant vectors, we propose a new heuristic for performing linear discriminant analysis in this high dimension, low sample size setting, based on projection onto the null space of $\boldsymbol{W}$, while simultaneously performing feature selection using sparsity inducing penalties.

The primary contributions of this paper are twofold. First, we propose a new heuristic for penalized classification based on ℓ_1-penalized zero-variance discriminant analysis; we provide a brief overview of zero-variance discriminant analysis in Section 2.1. The use of an ℓ_1-norm regularization term (or other sparsity inducing penalty) encourages sparse loading vectors for computing the low-dimensional representation, allowing further improvement in computational efficiency and interpretability. This approach in itself is not novel; ℓ_1-regularization and similar techniques have long been used in the statistics, machine learning,

and signal processing communities to induce sparse solutions, most notably in the LASSO [42] and compressed sensing [11, 34] regimes. However, our combination of ℓ_1-regularization and zero-variance discriminant analysis is novel and, based on experimental results summarized in Section 5, appears to improve on the current state of the art in terms of both computational efficiency and quality of discriminant vectors; a brief review of ℓ_1-regularization for high dimensional LDA may be found in Section 2.2.

Second, we propose a new algorithm for obtaining approximate solutions of penalized eigenproblems of the form (4.1), based on the alternating direction method of multipliers. Our algorithm finds an approximate solution of (4.1) by alternatingly maximizing each term of the objective function until convergence. Although the problem (4.1) has nonconcave objective in general, we will see that it is easy to maximize each term of the objective, either $\boldsymbol{x}^T\boldsymbol{B}\boldsymbol{x}$ or $-\|\boldsymbol{D}\boldsymbol{x}\|_1$, with the other fixed. We develop this algorithm in Section 3, as motivated by its use as a heuristic for penalized zero-variance discriminant analysis. Further, we show that this algorithm converges to a stationary point of (4.1) under certain assumptions on the matrices $\boldsymbol{W}$ and $\boldsymbol{D}$ in Section 3.2, and empirically test performance of the resulting classification heuristic in Section 5.

2 Linear Discriminant Analysis in the High Dimension, Low Sample Size Setting

2.1 Zero-Variance Linear Discriminant Analysis

As discussed in the previous section, the sample within-class covariance matrix $\boldsymbol{W}$ is singular in the high dimension, low sample size setting that occurs when $p > n$. Several solutions for this singularity problem have been proposed in the literature. One such proposed solution is to replace $\boldsymbol{W}$ in (1.3) with a positive definite approximation $\tilde{\boldsymbol{W}}$, e.g., the diagonal estimate $\tilde{\boldsymbol{W}} = \text{Diag}(\text{diag}(\boldsymbol{W}))$; see [5, 18, 22, 33, 48]. Here, $\text{Diag}(\boldsymbol{x})$ denotes the $p \times p$ diagonal matrix with diagonal entries given by $\boldsymbol{x} \in \mathbf{R}^p$ and $\text{diag}(\boldsymbol{X})$ denotes the vector in $\mathbf{R}^p$ with entries equal to those on the diagonal of $\boldsymbol{X} \in \mathbf{R}^{p\times p}$. After replacing the sample within-class covariance with a positive definite approximation $\tilde{\boldsymbol{W}}$, we obtain a set of discriminant vectors maximizing the modified LDA criterion by sequentially solving the generalized eigenproblems

$$\boldsymbol{w}_i = \underset{\boldsymbol{w}\in\mathbf{R}^p}{\arg\max} \left\{ \boldsymbol{w}^T\boldsymbol{B}\boldsymbol{w} : \boldsymbol{w}^T\tilde{\boldsymbol{W}}\boldsymbol{w} \leq 1,\ \boldsymbol{w}^T\tilde{\boldsymbol{W}}\boldsymbol{w}_j = 0\ \forall\, j = 1, \ldots, i-1 \right\} \tag{2.1}$$

as before. Unfortunately, it may often be difficult to obtain a good positive definite approximation of the population within-class scatter matrix. For example, the diagonal approxima-

tion ignores any correlation between features, while approximations based on perturbation of $\boldsymbol{W}$ may require some training to obtain a suitable choice of $\tilde{\boldsymbol{W}}$.

On the other hand, *zero-variance discriminant analysis* (ZVD), as proposed by Krzanowski et al. [33], embraces the singularity of $\boldsymbol{W}$ and seeks a set of discriminant vectors belonging to the null space of $\boldsymbol{W}$. If $\text{Null}(\boldsymbol{W}) \not\subseteq \text{Null}(\boldsymbol{B})$, we may obtain a nontrivial discriminant vector by solving the generalized eigenproblem

$$\max_{\boldsymbol{w}\in\mathbf{R}^p} \left\{ \boldsymbol{w}^T\boldsymbol{B}\boldsymbol{w} : \boldsymbol{W}\boldsymbol{w} = \boldsymbol{0},\ \boldsymbol{w}^T\boldsymbol{w} \leq 1 \right\}; \tag{2.2}$$

we may find the remaining nontrivial zero-variance discriminant vectors following a deflation process as before. That is, in ZVD we seek orthogonal directions $\boldsymbol{w}$ belonging to $\text{Null}(\boldsymbol{W})$ maximizing between-class scatter; because we are restricting our search to $\text{Null}(\boldsymbol{W})$, we seek orthogonal directions, not $\boldsymbol{W}$-conjugate directions as before. If the columns of $\boldsymbol{N} \in \mathbf{R}^{p\times d}$ form an orthonormal basis for $\text{Null}(\boldsymbol{W})$, then ZVD is equivalent to the eigenproblem

$$\max_{\boldsymbol{x}\in\mathbf{R}^d} \left\{ \boldsymbol{x}^T(\boldsymbol{N}^T\boldsymbol{B}\boldsymbol{N})\boldsymbol{x} : \boldsymbol{x}^T\boldsymbol{x} \leq 1 \right\}, \tag{2.3}$$

where d denotes the dimension of $\text{Null}(\boldsymbol{W})$. Clearly, the dimension of $\text{Null}(\boldsymbol{W}) \setminus \text{Null}(\boldsymbol{B})$ may be less than $k-1$. In this case, $\boldsymbol{N}^T\boldsymbol{B}\boldsymbol{N}$ has less than $k-1$ nontrivial eigenvectors; a full set of $k-1$ discriminant vectors can be obtained by searching for the remaining discriminant vectors in the complement of $\text{Null}(\boldsymbol{W})$ (see [19, pp. 8-9]). Alternately, reduced rank LDA could be performed using only the nontrivial discriminant vectors found in $\text{Null}(\boldsymbol{W})$.

2.2 Penalized Linear Discriminant Analysis

While ZVD and/or the use of a positive definite approximation of the within-class covariance matrix may solve the singularity problem in certain cases, particularly when the classes are separable in the subspace defined by the zero-variance discriminant vectors, the resulting discriminant vectors are typically difficult to interpret, especially when p is very large. Indeed, the method reduces to solving generalized eigenproblems and there is no reason to expect the obtained eigenvectors to possess any meaningful structure. Ideally, one would like to simultaneously perform feature selection and dimensionality reduction to obtain a set of discriminant vectors containing relatively few nonzero entries, or some other desired structure. In this case, one would be able to identify which features are truly meaningful for the task of the dimensionality reduction, while significantly improving computational efficiency through the use of sparse loading vectors.

Several recent articles have proposed variants of LDA involving ℓ_1-regularization as a surrogate for the restriction that the discriminant vectors be sparse in order to increase

interpretability of the obtained loading vectors. In [46], Witten and Tibshirani propose a penalized version of LDA where the kth discriminant vector is the solution of the optimization problem

$$\boldsymbol{w}_k = \underset{\boldsymbol{w}\in\mathbf{R}^p}{\arg\max} \left\{ \boldsymbol{w}^T\boldsymbol{B}\boldsymbol{w} - \rho(\boldsymbol{w}) : \boldsymbol{w}^T\tilde{\boldsymbol{W}}\boldsymbol{w} \leq 1,\ \boldsymbol{w}^T\tilde{\boldsymbol{W}}\boldsymbol{w}_i = 0\ \forall\, i \leq k-1 \right\}, \tag{2.4}$$

where $\rho : \mathbf{R}^p \to \mathbf{R}_+$ is either an ℓ_1-norm or fused LASSO penalty function, and $\tilde{\boldsymbol{W}}$ is the diagonal estimate of the within-class covariance $\tilde{\boldsymbol{W}} = \mathrm{Diag}(\mathrm{diag}(\boldsymbol{W}))$. The optimization problem (2.4) is nonconvex, because $\boldsymbol{B}$ is positive semidefinite, and cannot be solved as a generalized eigenproblem due to the presence of the regularization term $\rho(\boldsymbol{w})$. Consequently, it is unclear if it is possible to solve (2.4) efficiently. Witten and Tibshirani propose a minorization algorithm for approximately solving (2.4); the use of the diagonal estimate $\tilde{\boldsymbol{W}}$ is partially motivated by its facilitation of the use of soft thresholding when solving the subproblems arising in this minorization scheme.

Clemmensen et al. [14] consider an iterative method for penalized regression to obtain sparse discriminant vectors. Specifically, Clemmensen et al. apply an elastic net penalty [52] to the optimal scoring formulation of the LDA classification rule discussed in [25] as follows. Suppose that the first $\ell - 1$ discriminant vectors $\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{\ell-1}$ and scoring vectors $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \ldots, \boldsymbol{\theta}_{\ell-1}$ have been computed. Then the ℓth discriminant vector $\boldsymbol{w}_\ell$ and scoring vector $\boldsymbol{\theta}_\ell$ are the optimal solution pair of the problem

$$\begin{array}{ll} \min\limits_{\boldsymbol{w},\boldsymbol{\theta}} & \|\boldsymbol{Y}\boldsymbol{\theta} - \boldsymbol{X}\boldsymbol{w}\|^2 + \lambda_1 \boldsymbol{w}^T\boldsymbol{\Omega}\boldsymbol{w} + \lambda_2\|\boldsymbol{w}\|_1 \\ \text{s. t.} & \boldsymbol{\theta}^T\boldsymbol{Y}^T\boldsymbol{Y}\boldsymbol{\theta} = n,\ \boldsymbol{\theta}^T\boldsymbol{Y}^T\boldsymbol{Y}\boldsymbol{\theta}_\ell = 0\ \forall\, \ell < k. \end{array} \tag{2.5}$$

Here $\boldsymbol{Y}$ is the $n \times k$ partition matrix of the data set $\boldsymbol{X}$, i.e., Y_{ij} is the binary indicator variable for membership of the ith observation in the jth class, $\boldsymbol{\Omega}$ is a positive definite matrix chosen to ensure that $\boldsymbol{W} + \boldsymbol{\Omega}$ is positive definite and to encourage smoothness of the obtained discriminant vectors, and λ_1 and λ_2 are nonnegative tuning parameters controlling the ridge regression and ℓ_1-penalties, respectively. Clemmensen et al. propose the following iterative alternating direction method for solving (2.5). Suppose that we have the approximate solutions $\tilde{\boldsymbol{w}}_i$ and $\tilde{\boldsymbol{\theta}}_i$ of (2.5) at the ith step. These approximations are updated by first solving (2.5) for $\tilde{\boldsymbol{w}}_{i+1}$ with $\boldsymbol{\theta}$ fixed (and equal to $\tilde{\boldsymbol{\theta}}_i$); $\tilde{\boldsymbol{\theta}}_{i+1}$ is then updated by solving (2.5) with $\boldsymbol{w}$ fixed (and equal to $\tilde{\boldsymbol{w}}_{i+1}$). This process is repeated until the sequence of approximate solutions has converged or a maximum number of iterations has been performed. It is unclear if this algorithm is converging to a local minimizer because the criterion (2.5) is nonconvex; however, it can be shown that the solution of (2.4) is a stationary point of (2.5) under mild assumptions (see [46, Sect. 7.1]).

In addition to these heuristics, Cai and Liu (LPD) [10] propose a constrained ℓ_1-minimization for directly estimating the linear discriminant classifier for the two-class case, and establish

guarantees for both consistency and sparsity of the classifier. Moreover, several thresholding methods [24, 40, 43] for sparse LDA have also been proposed; a summary and numerical comparison of several of these cited methods can be found in [12].

2.3 Penalized Zero-Variance Linear Discriminant Analysis

Following the approach of [10, 14, 46], we propose a modification of zero-variance discriminant analysis, which we call SZVD, where ℓ_1-penalization is employed to induce sparse discriminant vectors. Specifically, we solve the problem

$$\max_{\boldsymbol{w}\in\mathbf{R}^p} \left\{ \tfrac{1}{2}\boldsymbol{w}^T\boldsymbol{B}\boldsymbol{w} - \gamma \textstyle\sum_{i=1}^p \sigma_i \left|(\boldsymbol{D}\boldsymbol{w})_i\right| : \boldsymbol{W}\boldsymbol{w} = \boldsymbol{0}, \boldsymbol{w}^T\boldsymbol{w} \leq 1 \right\} \tag{2.6}$$

to obtain the first discriminant vector; if the discriminant vectors $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_{k-1}$ have been identified, $\boldsymbol{w}_k$ can be found by appending $\{\boldsymbol{w}_1^T, \ldots, \boldsymbol{w}_{k-1}^T\}$ to the rows of $\boldsymbol{W}$ and solving (2.6) again. Here, $\boldsymbol{D} \in \mathcal{O}^p$ is an orthogonal matrix, and the ℓ_1-penalty acts as a convex surrogate for the cardinality of $\boldsymbol{w}$ with respect to the basis given by the columns of $\boldsymbol{D}$. The parameter $\boldsymbol{\sigma} \in \mathbf{R}^p$ is a scaling vector used to control emphasis of penalization; for example, the scaling vector $\boldsymbol{\sigma}$ may be taken to be the within-class standard deviations of the features $\boldsymbol{\sigma} = \sqrt{\operatorname{diag} \boldsymbol{W}}$ to ensure that a greater penalty is imposed on features that vary the most within each class. As before, letting the columns of $\boldsymbol{N} \in \mathbf{R}^{p\times d}$ form an orthonormal basis for Null($\boldsymbol{W}$) yields the equivalent formulation

$$\max_{\boldsymbol{x}\in\mathbf{R}^d} \left\{ \tfrac{1}{2}\boldsymbol{x}^T\boldsymbol{N}^T\boldsymbol{B}\boldsymbol{N}\boldsymbol{x} - \gamma \textstyle\sum_{i=1}^p \sigma_i \left|(\boldsymbol{D}\boldsymbol{N}\boldsymbol{x})_i\right| : \boldsymbol{x}^T\boldsymbol{x} \leq 1 \right\}. \tag{2.7}$$

Like (2.4), (2.7) is the maximization of a difference of convex functions over the unit ball; it is unknown if an efficient algorithm for solving (2.7) exists, although maximizing nonconcave functions is NP-hard in general. In the following section, we develop an algorithm to find stationary points of problem (2.7). Our algorithm is based on the *alternating direction method of multipliers* and we will use it as a heuristic for classification in Section 5.

3 SZVD: An Alternating Direction Method of Multipliers Approach for Penalized Zero-Variance Linear Discriminant Analysis

In this section, we develop an algorithm, called *SZVD*, for finding stationary points of the problem (2.7) based on the alternating direction method of multipliers; we will use these

stationary points as estimates of the global solution of (2.7) for the purpose of obtaining sparse discriminant vectors of a given data set. We develop our algorithm in Section 3.1 and establish convergence in Section 3.2.

3.1 The Alternating Direction Method of Multipliers

Given problems of the form $\min_{\boldsymbol{x},\boldsymbol{y}} \{f(\boldsymbol{x}) + g(\boldsymbol{y}) : \boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{y} = \mathbf{c}\}$, the *alternating direction method of multipliers* (ADMM) generates a sequence of iterates by alternately minimizing the augmented Lagrangian of the problem with respect to each primal decision variable, and then updating the dual variable using an approximate dual ascent; a recent survey on the ADMM and related methods can be found in [6].

To transform (2.7) to a form appropriate for the ADMM, we define an additional decision variable $\boldsymbol{y} \in \mathbf{R}^p$ by $\boldsymbol{y} = \boldsymbol{D}\boldsymbol{N}\boldsymbol{x}$. After this splitting of variables and replacing the maximization with an appropriate minimization, (2.7) is equivalent to

$$\min_{\boldsymbol{x}\in\mathbf{R}^d,\boldsymbol{y}\in\mathbf{R}^p} \left\{-\tfrac{1}{2}\boldsymbol{x}^T(\boldsymbol{N}^T\boldsymbol{B}\boldsymbol{N})\boldsymbol{x} + \gamma\textstyle\sum_{i=1}^p \sigma_i\,|y_i| : \boldsymbol{y}^T\boldsymbol{y} \le 1,\ \boldsymbol{D}\boldsymbol{N}\boldsymbol{x} = \boldsymbol{y}\right\}. \tag{3.1}$$

Letting $\boldsymbol{A} = \boldsymbol{N}^T\boldsymbol{B}\boldsymbol{N}$ and $\rho(\boldsymbol{y}) = \sum_{i=1}^p \sigma_i\,|y_i|$, we see that (3.1) is equivalent to

$$\min_{\boldsymbol{x}\in\mathbf{R}^d,\boldsymbol{y}\in\mathbf{R}^p} \left\{-\tfrac{1}{2}\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} + \gamma\rho(\boldsymbol{y}) : \boldsymbol{y}^T\boldsymbol{y} \le 1,\ \boldsymbol{D}\boldsymbol{N}\boldsymbol{x} - \boldsymbol{y} = \mathbf{0}\right\}. \tag{3.2}$$

This transformation has the additional benefit that $\rho(\boldsymbol{y})$ is separable in $\boldsymbol{y}$, while $\rho(\boldsymbol{D}\boldsymbol{N}\boldsymbol{x})$ is not separable in $\boldsymbol{x}$; this fact will play a significant role in our ADMM algorithm, as we will see shortly.

The problem (3.2) has augmented Lagrangian

$$L_\beta(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z}) = -\frac{1}{2}\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} + \gamma\rho(\boldsymbol{y}) + \delta_{B_p}(\boldsymbol{y}) + \boldsymbol{z}^T(\boldsymbol{D}\boldsymbol{N}\boldsymbol{x} - \boldsymbol{y}) + \frac{\beta}{2}\|\boldsymbol{D}\boldsymbol{N}\boldsymbol{x} - \boldsymbol{y}\|^2,$$

where $\beta > 0$ is a regularization parameter chosen so that L_β strictly convex in each of $\boldsymbol{x}$ and $\boldsymbol{y}$, and $\delta_{B_p} : \mathbf{R}^p \to \{0, +\infty\}$ is the indicator function of the ℓ_2-ball in $\mathbf{R}^p$, defined by $\delta_{B_p}(\boldsymbol{y}) = 0$ if $\boldsymbol{y}^T\boldsymbol{y} \le 1$ and equal to $+\infty$ otherwise. If we have the iterates $(\boldsymbol{x}^t, \boldsymbol{y}^t, \boldsymbol{z}^t)$ after t iterations, our algorithm obtains $(\boldsymbol{x}^{t+1}, \boldsymbol{y}^{t+1}, \boldsymbol{z}^{t+1})$ by sequentially minimizing L_β with respect to $\boldsymbol{y}$ and $\boldsymbol{x}$ and then updating $\boldsymbol{z}$ by an approximate dual ascent step:

$$\boldsymbol{y}^{t+1} = \operatorname*{arg\,min}_{\boldsymbol{y}\in\mathbf{R}^p} L_\beta\left(\boldsymbol{x}^t, \boldsymbol{y}, \boldsymbol{z}^t\right) \tag{3.3}$$

$$\boldsymbol{x}^{t+1} = \operatorname*{arg\,min}_{\boldsymbol{x}\in\mathbf{R}^p} L_\beta\left(\boldsymbol{x}, \boldsymbol{y}^{t+1}, \boldsymbol{z}^t\right) \tag{3.4}$$

$$\boldsymbol{z}^{t+1} = \boldsymbol{z}^t + \beta\left(\boldsymbol{D}\boldsymbol{N}\boldsymbol{x}^{t+1} - \boldsymbol{y}^{t+1}\right). \tag{3.5}$$

We next describe the solution of the subproblems (3.3) and (3.4).

We begin with (3.3). It is easy to see that (3.3) is equivalent to

$$\boldsymbol{y}^{t+1} = \operatorname*{arg\,min}_{\boldsymbol{y}\in\mathbf{R}^p} \left\{ \gamma\rho(\boldsymbol{y}) + \frac{\beta}{2}\boldsymbol{y}^T\boldsymbol{y} - \boldsymbol{y}^T(\beta\boldsymbol{D}\boldsymbol{N}\boldsymbol{x}^t + \boldsymbol{z}^t) : \boldsymbol{y}^T\boldsymbol{y} \leq 1 \right\}. \tag{3.6}$$

Applying the Karush-Kuhn-Tucker conditions [7, Section 5.5.3] to (3.6), we see that $\boldsymbol{y}^{t+1}$ must satisfy

$$\mathbf{0} \in \gamma\partial\rho(\boldsymbol{y}^{t+1}) + (\beta+\lambda)\boldsymbol{y}^{t+1} - (\beta\boldsymbol{D}\boldsymbol{N}\boldsymbol{x}^t + \boldsymbol{z}^t), \qquad \lambda((\boldsymbol{y}^{t+1})^T\boldsymbol{y}^{t+1} - 1) = 0 \tag{3.7}$$

for some Lagrange multiplier $\lambda \geq 0$. The stationarity condition and the form of the subdifferential of ρ (see [8, Section 3.4]) imply that there exists some $\boldsymbol{\phi} \in \mathbf{R}^p$, satisfying $\boldsymbol{\phi}^T\boldsymbol{y}^{t+1} = \rho(\boldsymbol{y}^{t+1})$ and $|\phi_i| \leq \sigma_i$ for all $i \in \{1,\ldots,p\}$, such that

$$0 = (\beta+\lambda)y_i^{t+1} + \gamma\phi_i - b_i,$$

where $\boldsymbol{b} = \beta\boldsymbol{D}\boldsymbol{N}\boldsymbol{x}^t + \boldsymbol{z}^t$, for each $i \in \{1,\ldots,p\}$. Rearranging and solving for y_i^{t+1} shows that

$$(\beta+\lambda)y_i^{t+1} = \operatorname{sign}(b_i)\cdot\max\{|b_i| - \gamma\sigma_i, 0\}$$

for all $i \in \{1,\ldots,p\}$. Letting $\boldsymbol{s}^{t+1} \in \mathbf{R}^p$ be the vector such that $s_i^{t+1} = \operatorname{sign}(b_i)\cdot\max\{|b_i| - \gamma\sigma_i, 0\}$ for all $i \in \{1,\ldots,p\}$ and applying the complementary slackness condition $\lambda((\boldsymbol{y}^{t+1})^T\boldsymbol{y}^{t+1} - 1) = 0$ shows that

$$\boldsymbol{y}^{t+1} = \frac{\boldsymbol{s}^{t+1}}{\beta + \max\{\mathbf{0}, \|\boldsymbol{s}^{t+1}\| - \beta\}}. \tag{3.8}$$

That is, we update $\boldsymbol{y}^{t+1}$ by applying soft thresholding to $\boldsymbol{b}$ with respect to $\gamma\boldsymbol{\sigma}$, and then normalizing the obtained solution if it has Euclidean norm greater than 1.

On the other hand, (3.4) is equivalent to

$$\boldsymbol{x}^{t+1} = \operatorname*{arg\,min}_{\boldsymbol{x}\in\mathbf{R}^d} \frac{1}{2}\boldsymbol{x}^T(\beta\boldsymbol{I} - \boldsymbol{A})\boldsymbol{x} + \boldsymbol{x}^T(\boldsymbol{D}\boldsymbol{N})^T(\boldsymbol{z}^t - \beta\boldsymbol{y}^{t+1}). \tag{3.9}$$

For sufficiently large choice of β, (3.9) is an unconstrained convex program. Finding the critical points of the objective of (3.9) shows that $\boldsymbol{x}^{t+1}$ is the solution of the linear system

$$(\beta\boldsymbol{I} - \boldsymbol{A})\boldsymbol{x}^{t+1} = (\boldsymbol{D}\boldsymbol{N})^T(\beta\boldsymbol{y}^{t+1} - \boldsymbol{z}^t). \tag{3.10}$$

The algorithm in Step 5 is stopped after t iterations if the primal and dual residuals satisfy

$$\|\boldsymbol{D}\boldsymbol{N}\boldsymbol{x}^t - \boldsymbol{y}^t\| \leq \epsilon_{abs}\cdot\sqrt{p} + \epsilon_{rel}\cdot\max\{\|\boldsymbol{x}^t\|, \|\boldsymbol{y}^t\|\}, \;\; \beta\|\boldsymbol{y}^t - \boldsymbol{y}^{t-1}\| \leq \epsilon_{abs}\cdot\sqrt{p} + \epsilon_{rel}\cdot\|\boldsymbol{y}^t\|$$

for desired absolute and relative error tolerances ϵ_{abs} and ϵ_{rel}; see [6, Sect. 3.3.1] for motivation for this choice of stopping criteria.

Algorithm 1 ADMM for Sparse zero-variance discriminant analysis (SZVD)

1: Given data $\boldsymbol{X}$. Compute sample between-class covariance $\boldsymbol{B}$ and within-class covariance $\boldsymbol{W}$.
2: Choose $k-1$ regularization parameters $\{\gamma_1, \gamma_2, \ldots, \gamma_{k-1}\}$, and sets of initial solutions $\{(\boldsymbol{x}^0, \boldsymbol{y}^0, \boldsymbol{z}^0)_i\}_{i=1}^{k-1}$. Set $i = 1$.
3: **repeat**
4: Compute basis $\boldsymbol{N}$ for the null space of $\boldsymbol{W}$.
5: Approximately solve (3.2) with regularization parameter $\gamma = \gamma_i$ using limit of iterates $\{(\boldsymbol{x}^t, \boldsymbol{y}^t, \boldsymbol{z}^t)\}_{t=0}^{\infty}$ generated by (3.8), (3.10), and (3.5) and the initial solution $(\boldsymbol{x}^0, \boldsymbol{y}^0, \boldsymbol{z}^0)_i$ to obtain $(\boldsymbol{x}^*, \boldsymbol{y}^*, \boldsymbol{z}^*)$. Let $\boldsymbol{w}_i = \boldsymbol{D}\boldsymbol{N}\boldsymbol{x}^* = \boldsymbol{y}^*$.
6: Append $\boldsymbol{w}_i$ to $\boldsymbol{W}$: $[\boldsymbol{W}; \boldsymbol{w}_i^T] \mapsto \boldsymbol{W}$.
7: Update $i = i + 1$.
8: **until** (all nontrivial zero-variance discriminant vectors $\{\boldsymbol{w}_i\}_{i=1}^{k-1}$ are computed.)

Our heuristic for identifying the (at most) $k-1$ penalized zero-variance discriminant vectors $\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{k-1}$ for a given data set $\boldsymbol{X}$ using the stationary points of (2.7) is summarized in Algorithm 1. The computation of $\boldsymbol{N}$ in Step 4 uses a QR decomposition of the within-class scatter matrix $\boldsymbol{W}$, computation of which requires $\mathcal{O}(p^3)$ operations. The ADMM Step 5 requires a preprocessing step consisting of Cholesky factorization of the coefficient matrix in (3.10) (at a cost of $\mathcal{O}(d^3)$ operations) and iterations containing a constant number of operations costing no more than matrix-matrix-vector multiplications of the form $\boldsymbol{D}\boldsymbol{N}\boldsymbol{v}$ (each requiring $\mathcal{O}(p^2)$ operations). It should be noted that the cost of this update can be significantly improved by exploiting the structure of $\boldsymbol{A}$ in some special cases. For example, if $k = 2$, we have the decomposition $\boldsymbol{A} = \boldsymbol{v}\boldsymbol{v}^T$, where $\boldsymbol{v} = \boldsymbol{N}^T(\hat{\boldsymbol{\mu}}_{\boldsymbol{1}} - \hat{\boldsymbol{\mu}}_{\boldsymbol{2}})$. Applying the Sherman-Morrison-Woodbury formula [23, Equation (2.1.4)] to (3.10) shows that

$$\boldsymbol{x}^{t+1} = \frac{1}{\beta}\left((\boldsymbol{D}\boldsymbol{N})^T(\beta\boldsymbol{y}^{t+1} - \boldsymbol{z}^t) - \left(\frac{(\beta\boldsymbol{y}^{t+1} - \boldsymbol{z}^t)^T\boldsymbol{D}\boldsymbol{N}\boldsymbol{v}}{\beta - \boldsymbol{v}^T\boldsymbol{v}}\right)\boldsymbol{v}\right),$$

which can be computed using $O(p^2)$ floating point operations (dominated by the matrix-vector multiplication $(\boldsymbol{D}\boldsymbol{N})^T(\beta\boldsymbol{y}^{t+1} - \boldsymbol{z}^t)$). That is, if $k = 2$, we use the available factorization $\boldsymbol{A} = \boldsymbol{v}\boldsymbol{v}^T$ to avoid the expensive Cholesky factorization step in Step 5. In either case, the algorithm has a total time complexity of $O((k-1)p^3) + O(\#\text{its} \cdot p^2)$, where #its denotes the total number of iterations required by Step 5.

3.2 Convergence Analysis

It is known that the ADMM converges to the optimal solution of

$$\min_{\boldsymbol{x}\in\mathbf{R}^{n_1},\boldsymbol{y}\in\mathbf{R}^{n_2}} \{f(\boldsymbol{x}) + g(\boldsymbol{y}) : \boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{y} = \boldsymbol{c}\}$$

if both f and g are convex functions (see [20, Theorem 8], [6, Section 3.2], and [27]) under mild assumptions on f, g, $\boldsymbol{A}$, and $\boldsymbol{B}$. However, general convergence results for minimizing nonconvex separable functions, such as the objective of (3.2), are unknown. Recently the authors of [28] have shown that the ADMM converges for certain nonconvex consensus and sharing problems. We note that the problems considered here are not covered by those considered in [28]; in particular, the constraints $\boldsymbol{D}\boldsymbol{N}\boldsymbol{x} = \boldsymbol{y}$ cannot be dealt with the analysis provided in [28].

In this section, we establish that, under certain assumptions on the within-class covariance matrix $\boldsymbol{W}$ and the dictionary matrix $\boldsymbol{D}$, the ADMM algorithm described by (3.3), (3.4), and (3.5) converges to a stationary point of (3.2). Let us define a new matrix $\boldsymbol{M} = \boldsymbol{D}\boldsymbol{N}$. Clearly the columns of $\boldsymbol{M}$ are also orthogonal, as we have $\boldsymbol{M}^T\boldsymbol{M} = \boldsymbol{N}^T\boldsymbol{D}^T\boldsymbol{D}\boldsymbol{N} = \boldsymbol{I}$. We have the following theorem.

Theorem 3.1 *Suppose that the columns of the matrix $[\boldsymbol{M}, \boldsymbol{C}] \in \mathcal{O}^p$ form an orthonormal basis for $\mathbf{R}^p$. Suppose further that the sequence of iterates $(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)$ generated by (3.3), (3.4), and (3.5) satisfies*

$$\boldsymbol{C}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k) = \boldsymbol{0} \tag{3.11}$$

for all k and the parameter parameter β satisfies

$$\beta > \|\boldsymbol{A}\| \left(\frac{\lambda_0^2 + 2}{\lambda_0^2}\right), \tag{3.12}$$

where $\|\boldsymbol{A}\|$ denotes the largest singular value of $\boldsymbol{A} = \boldsymbol{N}^T\boldsymbol{B}\boldsymbol{N}$, and λ_0 denotes the smallest nonzero eigenvalue of the matrix $\boldsymbol{M}\boldsymbol{M}^T$. Then $\left\{(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)\right\}_{k=0}^{\infty}$ converges to a stationary point of (3.2).

Although the assumption (3.11) that the successive difference of the multipliers belong to the null space of $\boldsymbol{M}$ may seem unrealistically restrictive, it is satisfied for two special classes of problems. We have the following corollary.

Corollary 3.1 *Suppose that $[\boldsymbol{M}, \boldsymbol{C}]$ forms the standard Euclidean basis for $\mathbf{R}^p$ and $\boldsymbol{z}^0 = \boldsymbol{M}\boldsymbol{a}^0$ for some vector $\boldsymbol{a}^0 \in \mathbf{R}^p$ with bounded norm. Then (3.11) is satisfied for all k and the sequence $\left\{(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)\right\}_{k=0}^{\infty}$ generated by (3.3), (3.4), and (3.5) converges to a stationary point of (3.2) if β satisfies (3.12).*

Proof: If $[\boldsymbol{M}, \boldsymbol{C}]$ forms the standard basis, then we can write $\boldsymbol{y} \in \mathbf{R}^p$ as $\boldsymbol{y} = \boldsymbol{M}\boldsymbol{c} + \boldsymbol{C}\boldsymbol{d}$ for some $\boldsymbol{c}$ and $\boldsymbol{d}$ with appropriate dimensions. It follows that we may decompose the subdifferential of the ℓ_1-norm at any $\boldsymbol{y} \in \mathbf{R}^p$ as

$$\partial\rho(\boldsymbol{y}) = \partial\rho\left(\boldsymbol{M}\boldsymbol{c} + \boldsymbol{C}\boldsymbol{d}\right) = \partial\rho\left(\boldsymbol{M}\boldsymbol{c}\right) + \partial\rho\left(\boldsymbol{C}\boldsymbol{d}\right) = \boldsymbol{M}\partial\rho\left(\boldsymbol{c}\right) + \boldsymbol{C}\partial\rho\left(\boldsymbol{d}\right) \tag{3.13}$$

by the fact that $\rho \circ \boldsymbol{M}$ and $\rho \circ \boldsymbol{C}$ are separable functions of $\boldsymbol{y}$. Substituting into the gradient condition of (3.7) we have

$$\begin{aligned} \mathbf{0} &\in \gamma\partial\rho(\boldsymbol{M}\boldsymbol{c}^1) + (\beta + \lambda)\boldsymbol{M}\boldsymbol{c}^1 - (\beta\boldsymbol{M}\boldsymbol{x}^0 + \boldsymbol{z}^0) \\ \mathbf{0} &\in \gamma\partial\rho(\boldsymbol{C}\boldsymbol{d}^1) + (\beta + \lambda)\boldsymbol{C}\boldsymbol{d}^1. \end{aligned}$$

Therefore, we must have $(\beta + \lambda)\boldsymbol{C}\boldsymbol{d}^1 \in -\gamma\partial\rho(\boldsymbol{C}\boldsymbol{d}^1)$, which implies that $\boldsymbol{C}\boldsymbol{d}^1 = \mathbf{0}$ by the structure of the subdifferential of the ℓ_1-norm. Extending this argument inductively shows that $\boldsymbol{C}\boldsymbol{d}^k = \mathbf{0}$ for all k, or equivalently, $\boldsymbol{C}^T\boldsymbol{y}^k = \mathbf{0}$ for all k. Substituting into (3.5) shows that

$$\boldsymbol{C}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k) = \beta\boldsymbol{C}^T(\boldsymbol{M}\boldsymbol{x}^{k+1} - \boldsymbol{y}^{k+1}) = -\beta\boldsymbol{C}^T(\boldsymbol{y}^{k+1}) = \mathbf{0}$$

because the columns of $\boldsymbol{C}$ are orthogonal to those of $\boldsymbol{M}$. This completes the proof. ■

Similarly, if $\boldsymbol{N}$ has both full row and column rank, as it would if $\boldsymbol{W} = \mathbf{0}$, then our ADMM algorithm converges. In particular, the algorithm converges when applied to (4.3) to identify the leading sparse principal component.

Corollary 3.2 *Suppose that $\boldsymbol{N}$ forms a basis for $\mathbf{R}^p$. Then* (3.11) *is satisfied for all k and the sequence $(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)$ generated by* (3.3), (3.4), *and* (3.5) *converges to a stationary point of* (3.2) *if β satisfies* (3.12).

Proof: If $\boldsymbol{N}$ forms a basis for $\mathbf{R}^p$, then $\boldsymbol{M}$ also forms a basis of $\mathbf{R}^p$, so its null space is spanned by $\boldsymbol{C} = \mathbf{0}$. Clearly $\boldsymbol{C}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k) = \mathbf{0}$ in this case. ■

The remainder of this section consists of a proof of Theorem 3.1. To establish Theorem 3.1, we will show that the value of the augmented Lagrangian of (3.2) decreases each iteration and the sequence of augmented Lagrangian values is bounded below. We will then exploit the fact that sequence of augmented Lagrangian values is convergent to show that the sequence of ADMM iterates is convergent. We conclude by establishing that a limit point of the sequence $\{(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)\}_{k=0}^{\infty}$ must be a stationary point of the original problem (3.2). We begin with the following lemma, which establishes that the augmented Lagrangian is decreasing provided the hypothesis of Theorem 3.1 is satisfied.

Lemma 3.1 *Suppose that* $\boldsymbol{C}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k) = \boldsymbol{0}$ *for all* k *and* $\beta > (\lambda_0^2 + 2)\|\boldsymbol{A}\|/\lambda_0^2$. *Then*

$$\begin{aligned} &L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^{k+1}) - L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k) \\ &\quad \leq -\frac{\beta}{2}\|\boldsymbol{y}^{k+1} - \boldsymbol{y}^k\|^2 - \frac{1}{2}\left(\beta - \|\boldsymbol{A}\| - \frac{2\|\boldsymbol{A}\|^2}{\beta\lambda_0}\right)\|\boldsymbol{x}^{k+1} - \boldsymbol{x}^k\|^2 \end{aligned} \tag{3.14}$$

and the right-hand side of (3.14) *is strictly negative if* $\boldsymbol{x}^{k+1} \neq \boldsymbol{x}^k$ *or* $\boldsymbol{y}^{k+1} \neq \boldsymbol{y}^k$.

Proof: We will obtain the necessary bound on the improvement in Lagrangian value given by $L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^{k+1}) - L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)$ by decomposing the difference as

$$\begin{aligned} &L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^{k+1}) - L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k) \\ &\quad = \left(L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^{k+1}) - L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^k)\right) \\ &\quad + \left(L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^k) - L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)\right) \end{aligned}$$

and bounding each summand in parentheses separately. We begin with the first summand $L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^{k+1}) - L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^k)$. Recall that $\boldsymbol{x}^{k+1}$ satisfies (cf. (3.10))

$$(\beta\boldsymbol{I} - \boldsymbol{A})\boldsymbol{x}^{k+1} = \boldsymbol{M}^T\left(\beta\boldsymbol{y}^{k+1} - \boldsymbol{z}^k\right).$$

Multiplying (3.5) by $\boldsymbol{M}^T$, using the fact that $\boldsymbol{M}^T\boldsymbol{M} = \boldsymbol{I}$, and substituting the formula above yields

$$\boldsymbol{M}^T\boldsymbol{z}^{k+1} = \boldsymbol{M}^T\boldsymbol{z}^k + \beta\boldsymbol{x}^{k+1} - \beta\boldsymbol{M}^T\boldsymbol{y}^{k+1} = \beta\boldsymbol{x}^{k+1} - \boldsymbol{M}^T(\beta\boldsymbol{y}^{k+1} - \boldsymbol{z}^k) = \boldsymbol{A}\boldsymbol{x}^{k+1}.$$

This implies that

$$\|\boldsymbol{M}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k)\| = \|\boldsymbol{A}(\boldsymbol{x}^{k+1} - \boldsymbol{x}^k)\| \leq \|\boldsymbol{A}\|\|\boldsymbol{x}^{k+1} - \boldsymbol{x}^k\|.$$

Applying the assumption (3.11), we have

$$\lambda_0\|\boldsymbol{z}^{k+1} - \boldsymbol{z}^k\| \leq \|\boldsymbol{M}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k)\| \leq \|\boldsymbol{A}\|\|\boldsymbol{x}^{k+1} - \boldsymbol{x}^k\|.$$

It follows immediately that

$$\begin{aligned} &L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^{k+1}) - L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^k) \\ &\quad = (\boldsymbol{z}^{k+1} - \boldsymbol{z}^k)^T(\boldsymbol{M}\boldsymbol{x}^{k+1} - \boldsymbol{y}^{k+1}) \\ &\quad = \frac{1}{\beta}\|\boldsymbol{z}^{k+1} - \boldsymbol{z}^k\|^2 \leq \frac{\|\boldsymbol{A}\|^2}{\beta\lambda_0^2}\|\boldsymbol{x}^{k+1} - \boldsymbol{x}^k\|^2. \end{aligned} \tag{3.15}$$

It remains to derive the necessary bound on $L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^k) - L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)$.

To do so, note that subproblem (3.9) is strongly convex with modulus $(\beta - \|\boldsymbol{A}\|)/2$. Let $f(\boldsymbol{x}) = -\frac{1}{2}\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} + \frac{\beta}{2}\|\boldsymbol{M}\boldsymbol{x} - \boldsymbol{y}^{k+1} + \boldsymbol{z}^k/\beta\|^2$. Then

$$\begin{aligned} L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^k) - L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^{k+1}, \boldsymbol{z}^k) &= f(\boldsymbol{x}^{k+1}) - f(\boldsymbol{x}^k) \\ &\leq -\nabla f(\boldsymbol{x}^{k+1})^T(\boldsymbol{x}^k - \boldsymbol{x}^{k+1}) - \left(\frac{\beta - \|\boldsymbol{A}\|}{2}\right)\|\boldsymbol{x}^{k+1} - \boldsymbol{x}^k\|^2 \\ &\leq -\left(\frac{\beta - \|\boldsymbol{A}\|}{2}\right)\|\boldsymbol{x}^{k+1} - \boldsymbol{x}^k\|^2 \end{aligned} \tag{3.16}$$

by the fact that $\boldsymbol{x}^{k+1}$ is a minimizer of f and, consequently, $\nabla f(\boldsymbol{x}^{k+1}) = 0$. Note that (3.15) and (3.16) in tandem imply that

$$L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^{k+1}) - L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^{k+1}, \boldsymbol{z}^k) \leq \frac{1}{2}\left(\|\boldsymbol{A}\| - \beta + \frac{2\|\boldsymbol{A}\|^2}{\beta\lambda_0^2}\right)\|\boldsymbol{x}^{k+1} - \boldsymbol{x}^k\|^2,$$

which is strictly negative if $\boldsymbol{x}^{k+1} \neq \boldsymbol{x}^k$, by the assumption that $\beta > (\lambda_0^2 + 2)\|\boldsymbol{A}\|/\lambda_0^2$.

Similarly, the subproblem (3.6) is strongly convex with modulus $\beta/2$. Let

$$g(\boldsymbol{y}) = \frac{\beta}{2}\boldsymbol{y}^T\boldsymbol{y} - \boldsymbol{y}^T(\beta\boldsymbol{M}\boldsymbol{x}^t + \boldsymbol{z}^t),$$

then from the optimality condition of problem (3.6), there must exist $\xi^{t+1} \in \partial\gamma\rho(\boldsymbol{y}^{k+1})$

$$\left(\xi^{k+1} + \nabla g(\boldsymbol{y}^{k+1})\right)^T(\boldsymbol{y} - \boldsymbol{y}^{k+1}) \geq 0, \quad \forall\ \boldsymbol{y} \in \mathbf{R}^p. \tag{3.17}$$

Then we have

$$\begin{aligned} &L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^{k+1}, \boldsymbol{z}^k) - L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k) \\ &\quad = \gamma\rho(\boldsymbol{y}^{k+1}) - \gamma\rho(\boldsymbol{y}^k) + g(\boldsymbol{y}^{k+1}) - g(\boldsymbol{y}^k) \\ &\quad \leq -\left(\xi^{k+1} + g(\boldsymbol{x}^{k+1})\right)^T(\boldsymbol{y}^k - \boldsymbol{y}^{k+1}) - \frac{\beta}{2}\|\boldsymbol{y}^{k+1} - \boldsymbol{y}^k\|^2 \\ &\quad \leq -\frac{\beta}{2}\|\boldsymbol{y}^{k+1} - \boldsymbol{y}^k\|^2. \end{aligned} \tag{3.18}$$

Combining (3.15), (3.16), and (3.18) gives the desired bound on the decrease of L. ■

Having established sufficient decrease of the augmented Lagrangian during each iteration, we next establish that the sequence of augmented Lagrangian values is bounded and, thus, convergent. We have the following lemma.

Lemma 3.2 *Suppose that $\boldsymbol{C}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k) = \boldsymbol{0}$ for all k and $\beta > (\lambda_0^2 + 2)\|\boldsymbol{A}\|/\lambda_0^2$. Then the sequence $\{L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)\}$ of augmented Lagrangian values is bounded. As a bounded monotonic sequence, $\{L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)\}$ is convergent.*

Proof: Note that the fact that $\boldsymbol{C}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k) = \boldsymbol{0}$ implies that $\boldsymbol{C}^T\boldsymbol{y}^{k+1} = \boldsymbol{0}$ for all k. Indeed, in this case

$$\boldsymbol{0} = \boldsymbol{C}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k) = \beta \boldsymbol{C}^T(\boldsymbol{M}\boldsymbol{x}^{k+1} - \boldsymbol{y}^{k+1}) = \beta \boldsymbol{C}^T\boldsymbol{y}^{k+1}$$

because the columns of $\boldsymbol{M}$ are orthogonal to those of $\boldsymbol{C}$. Thus, there exists $\{\boldsymbol{b}^k\}_{k=0}^{\infty} \in \mathbf{R}^d$, with $\|\boldsymbol{b}^k\| \leq 1$, such that $\boldsymbol{y}^k = \boldsymbol{M}\boldsymbol{b}^k$ for all k. In this case,

$$\begin{aligned}
L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k) &= -\frac{1}{2}(\boldsymbol{x}^k)^T\boldsymbol{A}\boldsymbol{x}^k + \gamma\rho(\boldsymbol{y}^k) + (\boldsymbol{z}^k)^T(\boldsymbol{M}\boldsymbol{x}^k - \boldsymbol{y}^k) + \frac{\beta}{2}\|\boldsymbol{M}\boldsymbol{x}^k - \boldsymbol{y}^k\|^2 \\
&= -\frac{1}{2}(\boldsymbol{x}^k)^T\boldsymbol{A}\boldsymbol{x}^k + \gamma\rho(\boldsymbol{y}^k) + (\boldsymbol{z}^k)^T\boldsymbol{M}(\boldsymbol{x}^k - \boldsymbol{b}^k) + \frac{\beta}{2}\|\boldsymbol{M}\boldsymbol{x}^k - \boldsymbol{y}^k\|^2 \\
&= -\frac{1}{2}(\boldsymbol{x}^k)^T\boldsymbol{A}\boldsymbol{x}^k + \gamma\rho(\boldsymbol{y}^k) + (\boldsymbol{A}\boldsymbol{x}^k)^T(\boldsymbol{x}^k - \boldsymbol{b}^k) + \frac{\beta}{2}\|\boldsymbol{M}\boldsymbol{x}^k - \boldsymbol{y}^k\|^2 \qquad (3.19) \\
&= -\frac{1}{2}(\boldsymbol{x}^k)^T\boldsymbol{A}\boldsymbol{x}^k + \gamma\rho(\boldsymbol{y}^k) + (\boldsymbol{A}\boldsymbol{x}^k)^T(\boldsymbol{x}^k - \boldsymbol{b}^k) + \frac{1}{2}(\boldsymbol{b}^k)^T\boldsymbol{A}\boldsymbol{b}^k \\
&\qquad - \frac{1}{2}(\boldsymbol{b}^k)^T\boldsymbol{A}\boldsymbol{b}^k + \frac{\beta}{2}\|\boldsymbol{M}\boldsymbol{x}^k - \boldsymbol{y}^k\|^2 \qquad (3.20) \\
&= \frac{1}{2}\Big(\|\boldsymbol{A}^{1/2}(\boldsymbol{x}^k - \boldsymbol{b}^k)\|^2 - (\boldsymbol{b}^k)^T\boldsymbol{A}\boldsymbol{b}^k + \beta\|\boldsymbol{M}\boldsymbol{x}^k - \boldsymbol{y}^k\|^2\Big) + \gamma\rho(\boldsymbol{y}^k)
\end{aligned}$$

where (3.19) follows from the identity $\boldsymbol{M}^T\boldsymbol{z}^k = \boldsymbol{A}\boldsymbol{x}^k$ and (3.20) follows from adding and subtracting $\frac{1}{2}(\boldsymbol{b}^k)^T\boldsymbol{A}\boldsymbol{b}^k$. Note that $\boldsymbol{A}^{1/2}$ is well-defined since $\boldsymbol{A}$ is a positive semidefinite matrix. Because both $\{\boldsymbol{b}^k\}$ and $\{\boldsymbol{y}^k\}$ are bounded, we may conclude that the sequence $\{L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k)\}$ is lower bounded. ■

As an immediate consequence of Lemma 3.2, we see that each of the sequences $\{\boldsymbol{x}^k\}$, $\{\boldsymbol{y}^k\}$, and $\{\boldsymbol{z}^k\}$ is convergent. Indeed, we have the following corollary.

Corollary 3.3 *Suppose that* $\boldsymbol{C}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k) = \boldsymbol{0}$ *for all* k *and* $\beta > (\lambda_0^2 + 2)\|\boldsymbol{A}\|/\lambda_0^2$. *Then* $\{\boldsymbol{x}^k\}$, $\{\boldsymbol{y}^k\}$, $\{\boldsymbol{z}^k\}$, *and* $\{\boldsymbol{M}\boldsymbol{x}^k - \boldsymbol{y}^k\}$ *are convergent, with*

$$\lim_{k\to\infty} \boldsymbol{M}\boldsymbol{x}^k - \boldsymbol{y}^k = \boldsymbol{0}.$$

Proof: The fact that $L_\beta(\boldsymbol{x}^{k+1}, \boldsymbol{y}^{k+1}, \boldsymbol{z}^{k+1}) - L_\beta(\boldsymbol{x}^k, \boldsymbol{y}^k, \boldsymbol{z}^k) \to \boldsymbol{0}$ and (3.14) imply that $\boldsymbol{x}^{k+1} - \boldsymbol{x}^k \to \boldsymbol{0}$ and $\boldsymbol{y}^{k+1} - \boldsymbol{y}^k \to \boldsymbol{0}$. The assumption that $\boldsymbol{C}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k) = \boldsymbol{0}$ and the identity $\boldsymbol{M}^T\boldsymbol{z}^k = \boldsymbol{A}\boldsymbol{x}^k$ implies that

$$\boldsymbol{M}^T(\boldsymbol{z}^{k+1} - \boldsymbol{z}^k) = \boldsymbol{A}(\boldsymbol{x}^{k+1} - \boldsymbol{x}^k) \to \boldsymbol{0}.$$

Thus, $\boldsymbol{z}^{k+1} - \boldsymbol{z}^k \to \boldsymbol{0}$ because the columns of $[\boldsymbol{M}, \boldsymbol{C}]$ form an orthonormal basis for $\mathbf{R}^p$, which further implies that the constraint violation satisfies $\boldsymbol{M}\boldsymbol{x}^k - \boldsymbol{y}^k = \boldsymbol{D}\boldsymbol{N}\boldsymbol{x}^k - \boldsymbol{y}^k \to \boldsymbol{0}$. This completes the proof ■

It remains to establish the following lemma, which states that any limit point of the sequence generated by (3.3), (3.4), and (3.5) is a stationary point of (3.2). This step is quite straightforward, but we include it here for completeness.

Lemma 3.3 *Let $\bar{\boldsymbol{x}}, \bar{\boldsymbol{y}}, \bar{\boldsymbol{z}}$ be limit points of the sequences $\{\boldsymbol{x}^k\}$, $\{\boldsymbol{y}^k\}$, and $\{\boldsymbol{z}^k\}$, respectively. Then*

$$\bar{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}\in\mathbf{R}^d} \left\{\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} + \bar{\boldsymbol{z}}^T(\boldsymbol{M}\boldsymbol{x} - \bar{\boldsymbol{y}})\right\} \tag{3.21}$$

$$\bar{\boldsymbol{y}} = \arg\min_{\boldsymbol{y}\in\mathbf{R}^p} \left\{\gamma\rho(\boldsymbol{y}) + \bar{\boldsymbol{z}}^T(\boldsymbol{M}\bar{\boldsymbol{x}} - \boldsymbol{y}) : \boldsymbol{y}^T\boldsymbol{y} \leq 1\right\} \tag{3.22}$$

$$\bar{\boldsymbol{y}} = \boldsymbol{M}\bar{\boldsymbol{x}}. \tag{3.23}$$

Therefore, $(\bar{\boldsymbol{x}}, \bar{\boldsymbol{y}}, \bar{\boldsymbol{z}})$ is a stationary point of $L_0(\bar{\boldsymbol{x}}, \bar{\boldsymbol{y}}, \bar{\boldsymbol{z}})$, i.e., $\mathbf{0} \in \partial L_0(\bar{\boldsymbol{x}}, \bar{\boldsymbol{y}}, \bar{\boldsymbol{z}})$.

Proof: That (3.23) holds is a consequence of the fact that $\boldsymbol{M}\boldsymbol{x}^k - \boldsymbol{y}^k \to \mathbf{0}$. The fact that $\mathbf{0} \in \partial L_0(\bar{\boldsymbol{x}}, \bar{\boldsymbol{y}}, \bar{\boldsymbol{z}})$ if (3.21) and (3.22) hold is an immediate consequence of the optimality conditions for the subproblems for $\boldsymbol{x}$ and $\boldsymbol{y}$ applied at $(\bar{\boldsymbol{x}}, \bar{\boldsymbol{y}}, \bar{\boldsymbol{z}})$. It remains to prove that (3.21) and (3.22) hold.

We begin with (3.21); (3.22) will follow by a similar argument. Fix k. Recall that $\boldsymbol{x}^{k+1}$ is a minimizer of the function $f(x) = -\frac{1}{2}\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} + \frac{\beta}{2}\|\boldsymbol{M}\boldsymbol{x} - \boldsymbol{y}^{k+1} + \boldsymbol{z}^k/\beta\|^2$. Therefore, $\boldsymbol{x}^{k+1}$ satisfies $\nabla f(\boldsymbol{x}^{k+1}) = 0$ for all $\boldsymbol{x} \in \mathbf{R}^d$. Evaluating the gradient of f at $\boldsymbol{x}^{k+1}$ shows that

$$\begin{aligned} 0 &\leq (\boldsymbol{x} - \boldsymbol{x}^{k+1})^T(-\boldsymbol{A}\boldsymbol{x}^{k+1} + \beta(\boldsymbol{x}^{k+1} - \boldsymbol{M}^T\boldsymbol{y}^{k+1}) + \boldsymbol{M}^T\boldsymbol{z}^k) \\ &= (\boldsymbol{x} - \boldsymbol{x}^{k+1})^T(-\boldsymbol{A}\boldsymbol{x}^{k+1} + \boldsymbol{M}^T\boldsymbol{z}^{k+1}) \end{aligned}$$

by (3.5). This implies that $\boldsymbol{x}^{k+1}$ is also a minimizer of $-\frac{1}{2}\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} + \boldsymbol{x}^T\boldsymbol{M}^T\boldsymbol{z}^{k+1}$. It follows that

$$-\frac{1}{2}\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} + (\boldsymbol{z}^{k+1})^T(\boldsymbol{M}\boldsymbol{x} - \boldsymbol{y}^{k+1}) \geq -\frac{1}{2}(\boldsymbol{x}^{k+1})^T\boldsymbol{A}\boldsymbol{x}^{k+1} + (\boldsymbol{z}^{k+1})^T(\boldsymbol{M}\boldsymbol{x}^{k+1} - \boldsymbol{y}^{k+1})$$

for all $\boldsymbol{x} \in \mathbf{R}^d$. Taking the limit as $k \to \infty$ shows that

$$-\frac{1}{2}\boldsymbol{x}^T\ \boldsymbol{A}\boldsymbol{x} + \bar{\boldsymbol{z}}^T(\boldsymbol{M}\boldsymbol{x} - \bar{\boldsymbol{y}}) \geq -\frac{1}{2}\bar{\boldsymbol{x}}^T\boldsymbol{A}\bar{\boldsymbol{x}}\ + \bar{\boldsymbol{z}}^T(\boldsymbol{M}\bar{\boldsymbol{x}} - \bar{\boldsymbol{y}})$$

for all $\boldsymbol{x} \in \mathbf{R}^d$, which establishes (3.21).

Let us look at the $\boldsymbol{y}$ update. Again let $g(\boldsymbol{y}) = \frac{\beta}{2}\boldsymbol{y}^T\boldsymbol{y} - \boldsymbol{y}^T(\beta\boldsymbol{M}\boldsymbol{x}^t + \boldsymbol{z}^t)$, then the optimality condition of this problem is (3.17). Writing this out explicitly, and noting the fact that for any two vectors $\xi_1 \in \partial\rho(\boldsymbol{y})$ and $\xi_2 \in \partial\rho(\boldsymbol{y}^{k+1})$ it must be true that $(\xi_1 - \xi_2)^T(\boldsymbol{y} - \boldsymbol{y}^{k+1}) \leq \rho(\boldsymbol{y}) - \rho(\boldsymbol{y}^{k+1})$, we have that every iterate $\boldsymbol{y}^{k+1}$ satisfies

$$\gamma(\rho(\boldsymbol{y}) - \rho(\boldsymbol{y}^{k+1})) + (\boldsymbol{y} - \boldsymbol{y}^{k+1})^T(\boldsymbol{z}^{k+1} + \beta\ \boldsymbol{M}(\boldsymbol{x}^{k+1} - \boldsymbol{x}^k)) \geq 0$$

for all $\boldsymbol{y} \in \mathbf{R}^p$. Taking the limit as $k \to \infty$ shows that

$$\gamma\rho(\boldsymbol{y}) + \bar{\boldsymbol{z}}^T(\boldsymbol{M}\bar{\boldsymbol{x}} - \boldsymbol{y}) \geq \gamma\rho(\bar{\boldsymbol{y}}) + \bar{\boldsymbol{z}}^T(\boldsymbol{M}\bar{\boldsymbol{x}} - \bar{\boldsymbol{y}})$$

as required. This completes the proof. ∎

4 Extension to Sparse Principal Component Analysis and Penalized Eigenproblems

The ADMM algorithm described by (3.3)-(3.5) may be applied to estimate the solution of *any* penalized eigenproblem of the form

$$\max_{\boldsymbol{x} \in \mathbf{R}^p} \left\{ -\tfrac{1}{2}\boldsymbol{x}^T\boldsymbol{B}\boldsymbol{x} + \gamma\|\boldsymbol{D}\boldsymbol{x}\|_1 : \boldsymbol{W}\boldsymbol{x} = \boldsymbol{0},\ \boldsymbol{x}^T\boldsymbol{x} \leq 1 \right\}. \tag{4.1}$$

That is, we make no assumption that $\boldsymbol{N} \in \mathbf{R}^{p \times d}$ forms a basis for the null space of the within-class covariance matrix of a labeled data set, but rather assume that the columns of $\boldsymbol{N}$ form an orthonormal basis corresponding to some subspace constraint.

The problem of identifying sparse solutions to eigenproblems has received significant attention, primarily in relation to sparse principal component analysis. In *principal component analysis* (PCA) [26, Section 14.5], one seeks a dimensionality reduction maximizing variance in the lower dimensional space. Specifically, the first k principal components are the k orthogonal directions $\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_k$ maximizing $\boldsymbol{w}^T\tilde{\boldsymbol{\Sigma}}\boldsymbol{w}$, where $\tilde{\boldsymbol{\Sigma}} \in \boldsymbol{S}_+^p$ is an approximation of the population covariance matrix (typically the sample covariance matrix); here $\boldsymbol{S}_+^p$ denotes the cone of $p \times p$ positive semidefinite matrices. Thus, principal component analysis reduces to identifying the k leading eigenvectors of the approximation of the covariance matrix given by $\tilde{\boldsymbol{\Sigma}}$. It is known that the sample covariance is a consistent estimator of the population covariance, i.e., the sample covariance matrix converges to the true population covariance matrix with probability 1 as the sample size n tends to infinity for fixed number of features p. However, when p is larger than n, as it is in the high dimension, low sample size setting, the sample covariance matrix may be a poor estimate of the population covariance; see [4,29,38]. One approach to addressing this issue, and to improve interpretability of the obtained loading vectors, is to require principal component vectors be sparse. Many different methods for this task have been proposed, typically involving ℓ_0 or ℓ_1-regularization, convex relaxation, thresholding, or some combination of all three; see [3, 15–17, 30, 31, 35, 37, 39, 47, 50, 51, 53] and the references within.

Our heuristic for LDA can be applied, with only minor modification, to a particular form of the sparse PCA problem. The leading principal component of a given data set can be

identified by solving the optimization problem

$$\boldsymbol{w}_1 = \operatorname*{arg\,max}_{\boldsymbol{w}\in\mathbf{R}^p} \left\{ \boldsymbol{w}^T\hat{\boldsymbol{\Sigma}}\boldsymbol{w} : \boldsymbol{w}^T\boldsymbol{w} \leq 1 \right\}, \tag{4.2}$$

where $\hat{\boldsymbol{\Sigma}} \in \boldsymbol{S}^p_+$ is the sample covariance matrix of the given data (after centering). A frequently used approach to simultaneously perform feature selection and principal component analysis is to restrict the obtained principal component to be s-sparse, with respect to the orthonormal basis $\boldsymbol{D} \in \mathcal{O}^p$, for some integer s:

$$\boldsymbol{w}_1 = \operatorname*{arg\,max}_{\boldsymbol{w}\in\mathbf{R}^p} \left\{ \boldsymbol{w}^T\hat{\boldsymbol{\Sigma}}\boldsymbol{w} : \boldsymbol{w}^T\boldsymbol{w} \leq 1,\ \|\boldsymbol{D}\boldsymbol{w}\|_0 \leq s \right\};$$

here $\|\boldsymbol{v}\|_0$ denotes the number of nonzero entries of the vector $\boldsymbol{v}$. Moving the cardinality constraint to the objective as a penalty and relaxing the ℓ_0-norm with the ℓ_1-norm yields the relaxation

$$\boldsymbol{w}_1 = \operatorname*{arg\,max}_{\boldsymbol{w}\in\mathbf{R}^p} \left\{ \boldsymbol{w}^T\hat{\boldsymbol{\Sigma}}\boldsymbol{w} - \gamma\|\boldsymbol{D}\boldsymbol{w}\|_1 : \boldsymbol{w}^T\boldsymbol{w} \leq 1 \right\}. \tag{4.3}$$

Clearly (4.3) is a special case of (2.7) with $\boldsymbol{B} = \hat{\boldsymbol{\Sigma}}$ and $\boldsymbol{W} = \boldsymbol{0}$ (or, equivalently, $\boldsymbol{N} = \boldsymbol{I}$), and $\boldsymbol{\sigma}$ equal to the all-ones vector of appropriate dimension. The relaxation (4.3) is very similar to the unified framework for sparse and functional principal component analysis of Allen [1], although we employ a different heuristic approach than the proximal operator-based alternating direction method considered in [1]. On the other hand, alternating direction methods and the method of multipliers have been used repeatedly in the literature to solve relaxations of the sparse PCA problem, most notably in [36], partially motivated by the ease of evaluating the proximal operator of the ℓ_1-norm.

We also note that our ADMM framework is immediately applicable to instances of (4.1) for *any* convex regularization function $\rho : \mathbf{R}^p \to \mathbf{R}$. For example, we may replace the ℓ_1-penalty $\sum \sigma_i|y_i|$ used above with the group lasso ℓ_1/ℓ_2 penalty [49], to induce discriminant vectors possessing desired group sparsity structure. However, we should remind the reader that the update step (3.3) requires evaluation of the proximal operator of $\rho + \delta_{B_p}$, which does not admit an analytic formula for general ρ.

5 Numerical Results

We performed a series of numerical experiments to compare our proposed algorithm (SZVD) with two recently proposed heuristics for penalized discriminant analysis, namely the PLDA [46] and SDA [14] methods discussed in Section 2.2 implemented as the `R` packages `penalizedLDA` [45] and `sparseLDA` [13] respectively.

In each experiment, we learn sets of $k-1$ discriminant vectors from given training data using each heuristic, and then test classification performance on a given test set. For each method, we apply validation to choose any regularization parameters. In particular, we choose the regularization parameters to minimize classification error if the discriminant vectors have at most a desired number of nonzero features; if the obtained discriminant vectors contain more nonzero entries than this tolerance, we use this number of nonzero entries as their validation score. We then use the regularization parameter γ corresponding to the minimum validation score, either classification error or number of nonzero features, to obtain our discriminant vectors. For each data set and $i = 1, \ldots, k-1$, this validation, applied to our ADMM heuristic (SZVD), selects the regularization parameter γ_i in (3.2) from a set of m evenly spaced values in the interval $[0, \tilde{\gamma}_i]$, where

$$\tilde{\gamma}_i := \frac{(\boldsymbol{w}_0)_i^T \boldsymbol{B} (\boldsymbol{w}_0)_i}{\rho((\boldsymbol{w}_0)_i)}$$

and $(\boldsymbol{w}_0)_i$ is the ith unpenalized zero-variance discriminant vector; this choice of $\tilde{\gamma}_i$ is made to ensure that the problem (3.2) has a nontrivial optimal solution by guaranteeing that at least one nontrivial solution with nonpositive objective value exists for each potential choice of γ_i. Similarly, in PLDA we perform validation on the tuning parameter λ controlling sparsity of the discriminant vectors in (2.4). Finally, SDA employs two tuning parameters, λ_1, which controls the ridge regression penalty, and λ_2 (`loads` in the R package), which controls the number of nonzero features; in each experiment we fix λ_1 and perform validation to choose λ_2. The choice of the weighting between classification performance and sparsity was chosen ad hoc to favor regularization parameters which yield moderately sparse solutions with very good misclassification rate; in experiments we observed that almost all classifiers found by SZVD had similar classification performance, regardless of regularization parameter, and we selected the parameter which gives the sparsest nontrivial solution via the validation step.

In all experiments, we use the dictionary matrix $\boldsymbol{D} = \boldsymbol{I}$, penalty weights $\boldsymbol{\sigma} = \text{diag}(\boldsymbol{W})$, regularization parameter $\beta = 2$, and stopping tolerances $tol_{abs} = tol_{rel} = 10^{-4}$ in SZVD. We initialized SZVD with the unpenalized zero-variance discriminant vectors $\boldsymbol{x}_1^0 = \boldsymbol{w}_1^0, \ldots, \boldsymbol{x}_{k-1}^0 = \boldsymbol{w}_{k-1}^0$ given by (2.3) and set $\boldsymbol{y}_i^0 = \boldsymbol{D}\boldsymbol{N}\boldsymbol{x}_i^0$ for all $i = 1, 2, \ldots, k-1$; consequently, we use the initial dual solution $\boldsymbol{z}_i^0 = \boldsymbol{0}$. Using this initialization, SZVD can be thought of as a refinement of the ZVD solution to induce sparsity. All features of discriminant vectors $\{\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{k-1}\}$ found using SZVD with magnitude less than 0.025 were rounded to 0, and all solutions containing trivial discriminant vectors were discarded. All experiments were performed in `R`; `R` and `Matlab` code for SZVD and for generating the synthetic data sets can be found at [2].

5.1 Simulated Data

Three sets of simulations were considered. For each $k \in \{2, 4\}$, we generate $25k$, $25k$, and $250k$ training, validation, and test observations in $\mathbf{R}^{500}$, respectively, as follows. In the first two, we sample 25, 25, and 250 Gaussian training, validation, and test observations from class C_i from the distribution $N(\boldsymbol{\mu}_i, \boldsymbol{\Sigma})$ for each $i \in \{1, 2, \ldots, k\}$, where the mean vector $\boldsymbol{\mu}_i$ is defined by

$$[\boldsymbol{\mu}_i]_j = \begin{cases} 0.7, & \text{if } 100(i-1)+1 \leq j \leq 100i \\ 0, & \text{otherwise} \end{cases}$$

and the covariance matrix $\boldsymbol{\Sigma}$ is chosen in one of two ways:

- In the first set of experiments, all features are correlated with

$$[\boldsymbol{\Sigma}_r]_{ab} = \begin{cases} 1, & \text{if } a = b \\ r, & \text{otherwise.} \end{cases}$$

  The experiment was repeated for each choice of $r \in \{0, 0.1, 0.5, 0.9\}$.

- In the second set of experiments, $\boldsymbol{\Sigma}$ is a block diagonal matrix with 100×100 diagonal blocks. For each (a, b) pair with a and b belonging to the same block, we have

$$[\boldsymbol{\Sigma}_\alpha]_{ab} = \alpha^{|a-b|}.$$

  We let $[\boldsymbol{\Sigma}_\alpha]_{ab} = 0$ for all remaining (a, b) pairs. The experiment was repeated for all choices of $\alpha \in \{0.1, 0.5, 0.9\}$.

In the third, we sample 25, 25, and 250 training, validation, and test observations from class C_i from the multivariate t distribution, where we use $\nu = 1$ degrees of freedom, and the mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ used in the underlying multivariate normal distribution function are chosen as in the first set of experiments above. To generate our random samples, we use the `R` package `mvtnorm` and the commands `rmvnorm` (for the Gaussians) and `rmvt` (for the t distribution).

For each (k, r) and (k, α) pair, we applied unpenalized zero-variance discriminant analysis (ZVD), our ADMM heuristic for penalized zero-variance discriminant analysis (SZVD), Witten and Tibshirani's penalized linear discriminant analysis with ℓ_1-norm and fused lasso penalties (PL1 and PFL), and the SDA algorithm of Clemmensen et al. (SDA) to obtain sets of $k-1$ discriminant vectors from the sampled training set. These discriminant vectors were then used to perform dimensionality reduction of the test data, and each observation in the test set was assigned to the class of the nearest projected training class centroid; an

identical process was applied to the validation data to train any regularization parameters. Both versions of PLDA chose the tuning parameter λ using 20 equally spaced values on the interval $[0, 0.15]$ by validation, while SDA used the ridge regression parameter $\lambda_1 = 10$ and chose the sparsity tuning parameter `loads` from the set

$$\{-500, -400, -300, -250, -200, -150, -120, -100, -80, -70, -60, -50\}$$

by validation. These ranges of λ and `loads` were chosen ad hoc to provided sets of potential discriminant vectors with a large range of classification performance and numbers of nonzero features for the sake of training the regularization parameters. The inner optimization of the SDA algorithm was stopped after a maximum of 5 iterations due to the high cost of each iteration. The desired sparsity level used in the validation process was at most 35% nonzero entries for each trial. This process was repeated 20 times for each (k, r) and (k, α) pair. All experiments were performed in `R` v.3.2.0 using one node of the University of Alabama's cluster RC2, containing two Intel Octa Core E5-2650 processors and 64GB of RAM. Tables 5.1, 5.2, and 5.3 report the average and standard deviation over all 20 trials for each set of simulated data of the number of misclassification errors, the number of nonzero features of discriminant vectors, and the time (in seconds) required to obtain each set of discriminant vectors per validation step for each of the five methods; the row headings *Err*, *Feat*, and *Time* correspond to the number of misclassification errors, the number of nonzero features, and the computation time, respectively.

5.2 Time-Series Data

We performed similar experiments for three data sets drawn from the UCR time series data repository [32], namely the *Coffee, OliveOil,* and *ECGFiveDays* data sets. The *ECGFiveDays* data set consists of 136-dimensional electrocardiogram measurements of a 67-year old male. Each observation corresponds to a measurement of the electrical signal of a single heartbeat of the patient. The data consists of two classes, 884 observations in total, corresponding to measurements taken on two dates, five days apart. We randomly divided the data into training, validation, and testing sets containing 25, 100, and 759 observations, respectively. We then applied each of our five methods to obtain discriminant vectors using each training and validation set pair and perform nearest centroid classification on the corresponding test set in the projected space. The tuning parameter λ in PLDA was selected from twenty equally spaced values in the interval $[0, 0.15]$, and we set $\lambda_1 = 0.001$ and chose the tuning parameter `loads` from the set

$$\{-500, -400, -300, -250, -200, -150, -120, -100, -80, -60, -50, -40, -30, -20, -10\}$$

Simulation 1		ZVD	SZVD	PL1	PFL	SDA
$k = 2$	*Err*	0.100 (0.447)	2.800 (4.697)	1.650 (2.323)	0.000 (0.000)	0.250 (0.444)
$r = 0$	*Feat*	490.350 (2.758)	136.850 (38.324)	121.750 (27.457)	186.300 (12.123)	149.350 (0.671)
	Time	0.494 (0.169)	0.468 (0.128)	0.017 (0.005)	0.024 (0.007)	1.212 (0.375)
$k = 2$	*Err*	0.100 (0.447)	5.100 (6.912)	9.100 (12.965)	26.350 (26.990)	2.300 (8.430)
$r = 0.1$	*Feat*	490.950 (2.645)	92.200 (27.156)	130.200 (25.622)	160.000 (39.901)	145.850 (15.749)
	Time	0.504 (0.177)	0.488 (0.200)	0.014 (0.004)	0.022 (0.006)	1.118 (0.276)
$k = 2$	*Err*	0.000 (0.000)	0.100 (0.447)	76.050 (59.797)	100.250 (63.519)	5.000 (11.671)
$r = 0.5$	*Feat*	489.300 (3.147)	111.300 (24.312)	146.600 (30.474)	143.350 (45.968)	149.650 (0.587)
	Time	0.512 (0.190)	0.473 (0.205)	0.017 (0.005)	0.029 (0.011)	1.236 (0.389)
$k = 2$	*Err*	0.000 (0.000)	0.000 (0.000)	140.600 (53.852)	156.550 (46.237)	29.950 (36.520)
$r = 0.9$	*Feat*	483.700 (3.585)	138.650 (7.286)	137.050 (42.908)	119.100 (41.573)	149.200 (1.005)
	Time	0.478 (0.156)	0.423 (0.146)	0.015 (0.003)	0.027 (0.008)	1.129 (0.294)
$k = 4$	*Err*	1.000 (1.026)	26.800 (32.413)	2.550 (2.305)	0.000 (0.000)	10.950 (16.340)
$r = 0$	*Feat*	1475.350 (5.851)	315.150 (39.833)	442.900 (131.772)	385.350 (19.258)	568.800 (424.873)
	Time	0.888 (0.314)	2.217 (0.697)	0.096 (0.030)	0.153 (0.043)	9.806 (4.014)
$k = 4$	*Err*	0.600 (0.883)	14.300 (7.168)	31.700 (30.418)	31.400 (29.564)	34.950 (37.870)
$r = 0.1$	*Feat*	1472.400 (5.051)	350.700 (84.639)	1176.500 (84.754)	1106.150 (315.328)	604.450 (398.613)
	Time	0.651 (0.166)	1.303 (0.075)	0.064 (0.002)	0.098 (0.003)	5.942 (1.409)
$k = 4$	*Err*	0.000 (0.000)	1.150 (1.814)	283.300 (99.537)	279.900 (100.690)	31.500 (30.645)
$r = 0.5$	*Feat*	1474.200 (4.641)	348.600 (57.135)	952.050 (471.463)	940.150 (455.762)	509.050 (381.087)
	Time	0.825 (0.263)	2.013 (0.653)	0.098 (0.032)	0.175 (0.067)	8.986 (2.921)
$k = 4$	*Err*	0.000 (0.000)	0.000 (0.000)	402.300 (85.345)	399.600 (87.247)	1.150 (3.233)
$r = 0.9$	*Feat*	1472.200 (5.616)	449.050 (118.077)	754.400 (466.590)	771.750 (485.399)	667.000 (419.571)
	Time	0.812 (0.276)	2.131 (0.768)	0.158 (0.069)	0.222 (0.088)	11.093 (4.220)

Table 5.1: Comparison of performance for synthetic data in $\mathbf{R}^{500}$ drawn from classes $C_1, \ldots, C_k \sim N(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_{\boldsymbol{r}})$ where $\boldsymbol{\Sigma}_{\boldsymbol{r}}$ is a 500×500 matrix with diagonal equal to 1 and all other entries equal to r. All values reported in the format "mean (standard deviation)". In all trials, $N_{train} = N_{val} = 25k$, $N_{test} = 250k$.

by validation when using SDA. As before, we stop the SDA inner optimization after 5 iterations. The desired sparsity level in the validation process was set to 25% nonzero entries for each training, validation, testing split. We repeated this process for 20 (training, validation, testing)-splits of the data and recorded the results in Table 5.4.

The *OliveOil* and *Coffee* data sets comprise 60 and 56 food spectrogram observations of different varieties of olive oil and coffee, respectively. Here, mass spectroscopy is applied to generate signals (spectra) corresponding to the molecular composition of samples of either coffee or olive oil. The goal is to distinguish between different varieties of olive oil and coffee

Simulation 2		ZVD	SZVD	PL1	PFL	SDA
$k = 2$	*Err*	0.200 (0.616)	8.200 (12.895)	3.250 (3.905)	0.000 (0.000)	0.400 (0.681)
$\alpha = 0.1$	*Feat*	490.300 (2.849)	103.350 (24.684)	115.350 (31.938)	179.100 (26.887)	149.200 (0.894)
	Time	0.297 (0.014)	0.291 (0.024)	0.010 (0.000)	0.014 (0.001)	0.731 (0.017)
$k = 2$	*Err*	10.100 (3.275)	28.500 (15.077)	9.450 (4.071)	6.550 (5.276)	9.000 (9.061)
$\alpha = 0.5$	*Feat*	491.350 (2.796)	102.000 (25.452)	130.950 (23.442)	170.800 (26.877)	135.550 (33.325)
	Time	0.302 (0.023)	0.291 (0.030)	0.010 (0.001)	0.015 (0.002)	0.737 (0.018)
$k = 2$	*Err*	212.500 (34.451)	223.700 (39.102)	89.500 (20.028)	91.300 (21.796)	91.300 (17.499)
$\alpha = 0.9$	*Feat*	490.550 (3.531)	106.350 (23.632)	137.250 (23.962)	143.650 (40.539)	99.100 (33.606)
	Time	0.331 (0.074)	0.418 (0.112)	0.012 (0.004)	0.018 (0.004)	0.789 (0.119)
$k = 4$	*Err*	1.850 (1.424)	28.600 (7.423)	3.900 (3.076)	0.650 (1.137)	10.800 (14.468)
$\alpha = 0.1$	*Feat*	1474.150 (4.793)	306.300 (21.957)	423.350 (82.933)	383.300 (19.834)	542.250 (371.578)
	Time	0.572 (0.046)	1.398 (0.137)	0.067 (0.006)	0.109 (0.022)	6.319 (1.385)
$k = 4$	*Err*	55.650 (12.287)	115.750 (22.856)	25.900 (12.397)	18.750 (8.410)	59.950 (43.584)
$\alpha = 0.5$	*Feat*	1476.100 (6.290)	310.400 (29.170)	738.050 (307.163)	721.300 (304.402)	580.100 (487.021)
	Time	0.596 (0.124)	1.451 (0.154)	0.067 (0.003)	0.106 (0.007)	7.052 (1.662)
$k = 4$	*Errs*	477.200 (28.957)	492.900 (28.306)	300.100 (18.379)	299.650 (18.088)	380.400 (39.487)
$\alpha = 0.9$	*Feat*	1474.200 (4.841)	402.700 (26.872)	1056.300 (337.691)	1075.400 (314.011)	256.650 (127.942)
	Time	0.582 (0.022)	2.045 (0.103)	0.068 (0.005)	0.122 (0.066)	6.679 (0.468)

Table 5.2: Comparison of performance for synthetic data in $\mathbf{R}^{500}$ drawn from classes $C_1, \ldots, C_k \sim N(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_{\boldsymbol{\alpha}})$ where $\boldsymbol{\Sigma}_{\boldsymbol{\alpha}}$ is a 500×500 diagonal block matrix with 100×100 diagonal blocks with (i, j) nonzero entry equal to $\alpha^{|i-j|}$. In all trials, $N_{train} = N_{val} = 25k$, $N_{test} = 250k$.

from these spectral signals. The *OliveOil* data set [41] consists of 570-dimensional spectrograms corresponding to samples of extra virgin olive oil from one of four countries (Greece, Italy, Portugal, or Spain); 286-dimensional spectrograms of either Arabica or Robusta variants of instant coffee compose the *Coffee* data set [9]. As before, we divide the *OliveOil* data into training, validation, and testing sets containing $(30, 10, 20)$ observations, respectively. We then applied each of the five approaches to learn a classification rule from the training and validation data and classify the given test data. We used the same range of tuning parameter λ as in the *ECGFiveDays* trials for each PLDA heuristic; we stopped the inner optimization step after 5 iterations, set $\lambda_1 = 0.1$, and used the same set of potential values of the tuning parameter `loads` as we did in *ECGFiveDays* trials for SDA. This process was repeated for 20 different data splits, and we then repeated the experiment for the *Coffee* data set using training, validation, test splits of size $(25, 10, 21)$. For both *Coffee* and *OliveOil*, the desired sparsity level in the validation process was set to 35% nonzero entries for each training, validation, testing split. The results of these trials are summarized in Table 5.4.

Simulation 3		ZVD	SZVD	PL1	PFL	SDA
$k = 2$	*Err*	106.500 (22.090)	147.500 (26.437)	231.400 (32.173)	219.050 (48.424)	178.900 (60.225)
$r = 0$	*Feat*	486.500 (3.348)	89.100 (19.040)	233.600 (79.962)	277.700 (86.270)	125.300 (37.480)
	Time	0.297 (0.020)	0.558 (0.096)	0.009 (0.005)	0.007 (0.001)	0.674 (0.027)
$k = 2$	*Err*	111.800 (28.298)	150.900 (31.692)	216.400 (41.140)	222.650 (36.152)	182.950 (61.908)
$r = 0.1$	*Feat*	486.600 (3.470)	99.300 (20.538)	227.150 (127.067)	295.600 (54.472)	139.150 (18.466)
	Time	0.289 (0.014)	0.539 (0.106)	0.007 (0.001)	0.007 (0.001)	0.676 (0.029)
$k = 2$	*Err*	76.200 (29.890)	140.200 (123.300)	224.850 (26.688)	223.000 (30.184)	181.750 (55.999)
$r = 0.5$	*Feat*	488.550 (3.605)	102.950 (41.945)	235.000 (124.283)	307.650 (99.222)	127.750 (32.598)
	Time	0.287 (0.010)	0.508 (0.122)	0.008 (0.001)	0.008 (0.001)	0.647 (0.018)
$k = 2$	*Err*	25.700 (8.832)	76.700 (140.006)	230.250 (22.178)	229.250 (24.244)	172.550 (86.696)
$r = 0.9$	*Feat*	484.700 (3.466)	169.450 (58.531)	350.600 (134.481)	364.300 (158.035)	140.200 (12.903)
	Time	0.284 (0.009)	0.364 (0.122)	0.007 (0.001)	0.011 (0.002)	0.630 (0.017)
$k = 4$	*Err*	206.700 (28.488)	254.650 (33.971)	689.550 (85.829)	677.750 (101.469)	632.800 (128.331)
$r = 0$	*Feat*	1469.950 (4.883)	483.300 (128.665)	610.900 (372.315)	698.900 (259.122)	580.600 (445.667)
	Time	0.539 (0.014)	1.916 (0.428)	0.044 (0.002)	0.031 (0.003)	5.660 (1.169)
$k = 4$	*Err*	203.050 (40.382)	249.950 (36.767)	700.850 (53.635)	686.950 (61.275)	583.150 (108.370)
$r = 0.1$	*Feat*	1468.900 (6.965)	467.400 (119.478)	712.900 (388.527)	873.750 (324.816)	539.850 (389.930)
	Time	0.538 (0.011)	1.842 (0.415)	0.043 (0.003)	0.032 (0.002)	5.495 (1.053)
$k = 4$	*Err*	131.450 (34.861)	204.900 (84.124)	696.600 (55.645)	691.400 (66.199)	608.450 (127.759)
$r = 0.5$	*Feat*	1470.350 (6.467)	534.950 (187.540)	671.200 (304.304)	780.850 (329.883)	570.950 (386.541)
	Time	0.661 (0.212)	1.910 (0.625)	0.047 (0.008)	0.045 (0.007)	5.889 (1.160)
$k = 4$	*Err*	50.700 (13.929)	189.950 (177.184)	711.000 (43.623)	703.600 (48.309)	487.100 (228.919)
$r = 0.9$	*Feat*	1472.450 (5.246)	644.450 (142.944)	751.750 (252.882)	878.350 (336.339)	849.800 (446.208)
	Time	0.691 (0.217)	1.868 (0.704)	0.045 (0.003)	0.058 (0.009)	6.444 (1.411)

Table 5.3: Comparison of performance for synthetic data in $\mathbf{R}^{500}$ drawn from multivariate t distribution as described above. All values reported in the format "mean (standard deviation)". In all trials, $N_{train} = N_{val} = 25k$, $N_{test} = 250k$.

5.3 Commentary

As can be seen from the experiments of the previous section, our proposed algorithm SZVD compares favorably to the current state of the art. When compared to the zero-variance discriminant, adding penalization in the form of a ℓ_1-penalty results in a modest degradation in classification performance, as may be expected. However, this penalization significantly increased sparsity of the obtained discriminant vectors from that of the zero-variance discriminants. Moreover, SZVD significantly outperforms both forms of PLDA in terms of classification error, and exhibits similar performance to SDA while using significantly fewer computing resources. In particular, our heuristic exhibits better classification performance,

Time-Series		ZVD	SZVD	PL1	PFL	SDA
OliveOil	*Err*	1.900 (1.373)	2.900 (1.714)	3.700 (1.625)	4.100 (1.861)	1.950 (1.432)
$p = 570$	*Feat*	1669.600 (6.723)	312.050 (75.076)	331.900 (156.725)	343.750 (252.559)	969.050 (358.902)
$k = 4$	*Time*	1.052 (0.151)	5.027 (0.799)	0.020 (0.005)	0.046 (0.006)	5.788 (0.841)
Coffee	*Err*	0.000 (0.000)	0.050 (0.224)	6.600 (2.371)	6.800 (2.462)	0.250 (0.639)
$p = 286$	*Feat*	281.950 (2.212)	44.250 (10.432)	123.700 (56.610)	129.250 (67.392)	55.400 (25.046)
$k = 2$	*Time*	0.098 (0.015)	0.269 (0.058)	0.009 (0.001)	0.013 (0.002)	0.371 (0.010)
ECG	*Err*	33.850 (18.253)	45.100 (21.506)	145.100 (26.989)	152.350 (28.027)	52.000 (24.921)
$p = 136$	*Feat*	134.850 (0.671)	32.200 (9.192)	34.200 (11.190)	38.300 (13.692)	26.500 (11.821)
$k = 2$	*Time*	0.027 (0.008)	0.068 (0.016)	0.009 (0.001)	0.014 (0.003)	0.211 (0.019)

Table 5.4: Comparison of performance for the *OliveOil, Coffee, and ECGFiveDays* data sets.

when compared to the three penalized LDA heuristics for all but the uncorrelated synthetic data sets in Simulation 1, and two of the three real data sets. Similarly, SZVD uses the fewest features in all of the four class trials in Simulation 1, and at least second fewest in all other trials. Finally, SZVD uses less computing resources than all but the ZVD and PLDA heuristics, which perform poorly in terms of number of features and classification, respectively. It should also be noted that, although we did not verify that the conditions ensuring convergence of our ADMM heuristic given by Theorem 3.1 are satisfied (and there is no reason to expect them to be), the ADMM heuristic converged in all trials after at most a few hundred iterations.

There were two notable exceptions. First, for uncorrelated data, i.e., the $r = 0$ case in Simulation 1, both variants of PLDA outperformed SZVD in terms of classification error. This is not surprising, as the implicit assumption that the data are uncorrelated made when using the diagonal approximation holds for this special case. However, the performance of PLDA degrades significantly as r is increased, while that of SZVD improves. Second, SZVD performs very poorly for highly correlated data in Simulation 2; roughly half of all test observations are misclassified in the $\alpha = 0.9$ trials. It should be noted that SZVD performs only marginally worse than unpenalized zero-variance discriminant analysis (ZVD) for $\alpha = 0.5$ and $\alpha = 0.9$; there is a significantly sharper decrease in classification performance from ZVD to SZVD when the data are weakly correlated ($r = 0$ and $r = 1$ in Simulation 1 and $\alpha = 0.1$ in Simulation 2). This suggests that the classes may not be linearly separable in the null space of the within-class covariance matrix in this case. It should also be noted that none of the heuristics perform well for these particular synthetic data sets, with PLDA and SDA misclassifying at least one third of test observations on average. Finally, none of the compared methods performed well, in terms of classification rate, when applied to the data sampled from multivariate t distributions. This is somewhat to be expected, as Fisher's

LDA is designed for symmetric distributions, such as the normal distribution, while the t distribution is asymmetric. However, it should be noted that ZVD and SZVD exhibited significantly better classification performance than PLDA and SDA.

The use of penalization, aside from encouraging sparsity of the discriminant vectors, also seems to increase interpretability of the discriminant vectors, although it should be noted that this claim is somewhat subjective. For example, the nonzero entries of the discriminant vector used to classify the *ECGFiveDays* data align with features where data in different classes appear to differ significantly (on average). Specifically, both the zero-variance and SZVD discriminant vectors closely follow the trajectory of the difference of the sample class-mean vectors $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2$. However, most of the entries of the SZVD discriminant vector corresponding to small magnitude entries of the zero-variance discriminant are set equal to zero; the remaining nonzero entries of the SZVD discriminant vector correspond to features where the two class mean vectors seem to differ the most significantly. This is most apparent when comparing the discriminant vectors to the class mean vectors of the data after centering and normalizing, although this phenomena is also weakly visible when comparing class means for the original data set. See Figure 5.1 for more details.

6 Conclusions

We have developed a novel heuristic for simultaneous feature selection and linear discriminant analysis based on penalized zero-variance discriminant analysis and the alternating direction method of multipliers. Our approach offers several advantages over the current state of the art. Most notably, our algorithm, SZVD, employs inexpensive iterations, relative to those of the heuristics employed in PLDA [46] and SDA [14]. One exception to this claim occurs when the diagonal estimate of $\boldsymbol{W}$ is used in PLDA; in this case, PLDA is faster than our algorithm (as seen in the experiments of Section 5). However, our algorithm appears to provide significantly better classification performance than this particular version of PLDA, likely due to the use, by SZVD, of information regarding correlation of features ignored by the diagonal estimate. Moreover, our algorithm can be shown to converge under some conditions on $\boldsymbol{W}$ and $\boldsymbol{D}$ (see Section 3.2). On the other hand, our approach is applicable, without modification, to the case when $k \geq 3$, unlike the linear programming approach of [10].

This work opens several interesting areas of future research. For example, we establish convergence of the ADMM when applied to a nonconvex optimization problem by the novel use of the augmented Lagrangian as a measure of progress in our convergence analysis. It would be extremely interesting to see if this approach and analysis can be extended to

other classes of nonconvex optimization problems. On the other hand, it would be useful to strengthen our analysis to obtain estimates of the convergence rate of our algorithm and, if possible, establish convergence of SZVD in the absence of the assumptions on $\boldsymbol{W}$ and $\boldsymbol{D}$ made in Section 3.2. Finally, future research might include consistency analysis of penalized zero-variance analysis to identify when our heuristic is able to correctly classify data and identify important features.

7 Acknowledgments

The work presented in this paper was partially carried out while B.P.W Ames was a postdoctoral fellow at the Institute for Mathematics and its Applications during the IMA's annual program on Mathematics of Information (supported by National Science Foundation (NSF) award DMS-0931945), and while B.P.W. Ames was a Von Karman instructor at the California Institute of Technology supported by Joel Tropp under Office of Naval Research (ONR) award N00014-11-1002. This research was also supported by University of Alabama Research Grant RG14678. We are grateful to Fadil Santosa, Krystal Taylor, Zhi-Quan Luo, Meisam Razaviyayn, and Line Clemmensen for their insights and helpful suggestions. Finally, we are grateful for the contributions of two anonymous reviewers whose suggestions have greatly improved this manuscript.

References

[1] G. I. Allen. Sparse and functional principal components analysis. *arXiv preprint arXiv:1309.2895*, 2013. 18

[2] B. Ames and M. Hong. SZVD: ADMM for sparse zero-variance discriminant analysis. `http://bpames.people.ua.edu/software.html`, 2014. 19

[3] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Structured sparsity through convex optimization. *Statistical Science*, 27(4):450–468, 2012. 17

[4] J. Baik and J. W. Silverstein. Eigenvalues of large sample covariance matrices of spiked population models. *Journal of Multivariate Analysis*, 97(6):1382–1408, 2006. 17

[5] P. Bickel and E. Levina. Some theory for Fisher's linear discriminant function, naive Bayes', and some alternatives when there are many more variables than observations. *Bernoulli*, 10(6):989–1010, 2004. 4

[6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011. 8, 9, 11

[7] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, UK, 2004. 9

[8] S. Boyd and L. Vandenberghe. Subgradients. *Lecture Notes for EE364b, Stanford University, Winter 2006-07*, 2008. Available from `http://see.stanford.edu/materials/lsocoee364b/01-subgradients_notes.pdf`. 9

[9] R. Briandet, E. K. Kemsley, and R. H. Wilson. Discrimination of Arabica and Robusta in instant coffee by Fourier transform infrared spectroscopy and chemometrics. *Journal of agricultural and food chemistry*, 44(1):170–174, 1996. 23

[10] T. Cai and W. Liu. A direct estimation approach to sparse linear discriminant analysis. *Journal of the American Statistical Association*, 106(496), 2011. 6, 7, 26

[11] E. J. Candès. Compressive sampling. In *Proceedings oh the International Congress of Mathematicians: Madrid, August 22-30, 2006: invited lectures*, pages 1433–1452, 2006. 4

[12] L. Clemmensen. *On discriminant analysis techniques and correlation structures in high dimensions*. Technical Report-2013. Technical University of Denmark, 2013. 7

[13] L. Clemmensen and contributions by Max Kuhn. *sparseLDA: Sparse Discriminant Analysis*, 2012. R package version 0.1-6. 18

[14] L. Clemmensen, T. Hastie, D. Witten, and B. Ersbøll. Sparse discriminant analysis. *Technometrics*, 53(4), 2011. 6, 7, 18, 26

[15] A. d'Aspremont, F. Bach, and L. El Ghaoui. Optimal solutions for sparse principal component analysis. *J. Mach. Learn. Res.*, 9:1269–1294, 2008. 17

[16] A. d'Aspremont, F. Bach, and L. El Ghaoui. Approximation bounds for sparse principal component analysis. *Mathematical Programming*, 148(1-2):89–110, 2014. 17

[17] A. d'Aspremont, L. El Ghaoui, M. I. Jordan, and G. R. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. *SIAM review*, 49(3):434–448, 2007. 17

[18] S. Dudoit, J. Fridlyand, and T. P. Speed. Comparison of discrimination methods for the classification of tumors using gene expression data. *Journal of the American statistical association*, 97(457):77–87, 2002. 4

[19] J. Duintjer Tebbens and P. Schlesinger. Improving implementation of linear discriminant analysis for the high dimension/small sample size problem. *Computational Statistics & Data Analysis*, 52(1):423–437, 2007. 5

[20] J. Eckstein and D. P. Bertsekas. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1-3):293–318, 1992. 11

[21] R. A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7(2):179–188, 1936. 2

[22] J. H. Friedman. Regularized discriminant analysis. *Journal of the American statistical association*, 84(405):165–175, 1989. 4

[23] G. H. Golub and C. F. Van Loan. *Matrix computations*. Johns Hopkins University Press, 1996. 10

[24] Y. Guo, T. Hastie, and R. Tibshirani. Regularized linear discriminant analysis and its application in microarrays. *Biostatistics*, 8(1):86–100, 2007. 7

[25] T. Hastie, A. Buja, and R. Tibshirani. Penalized discriminant analysis. *The Annals of Statistics*, pages 73–102, 1995. 6

[26] T. Hastie, R. Tibshirani, and J. J. H. Friedman. *The elements of statistical learning*. Springer New York, 2013. 2, 17

[27] M. Hong and Z.-Q. Luo. On the linear convergence of the alternating direction method of multipliers. *arXiv preprint arXiv:1208.3922*, 2012. 11

[28] M. Hong, Z.-Q. Luo, and M. Razaviyayn. Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems. *arXiv preprint arXiv:1410.1390*, 2012. 11

[29] I. M. Johnstone and A. Y. Lu. Sparse principal components analysis. *Unpublished manuscript*, 2004. 17

[30] I. T. Jolliffe, N. T. Trendafilov, and M. Uddin. A modified principal component technique based on the lasso. *Journal of Computational and Graphical Statistics*, 12(3):531–547, 2003. 17

[31] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre. Generalized power method for sparse principal component analysis. *The Journal of Machine Learning Research*, 11:517–553, 2010. 17

[32] E. Keogh, X. Xi, L. Wei, and C. A. Ratanamahatana. The UCR time series classification/clustering homepage, 2006. 21

[33] W. Krzanowski, P. Jonathan, W. McCarthy, and M. Thomas. Discriminant analysis with singular covariance matrices: methods and applications to spectroscopic data. *Applied statistics*, pages 101–115, 1995. 4, 5

[34] G. Kutyniok. Compressed sensing: Theory and applications. *CoRR, vol. abs/1203.3815*, 2012. 4

[35] R. Luss and M. Teboulle. Convex approximations to sparse PCA via lagrangian duality. *Operations Research Letters*, 39(1):57–61, 2011. 17

[36] S. Ma. Alternating direction method of multipliers for sparse principal component analysis. *Journal of the Operations Research Society of China*, 1(2):253–274, 2013. 18

[37] D. S. Papailiopoulos, A. G. Dimakis, and S. Korokythakis. Sparse PCA through low-rank approximations. *arXiv preprint arXiv:1303.0551*, 2013. 17

[38] D. Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17(4):1617, 2007. 17

[39] P. Richtárik, M. Takáč, and S. D. Ahipaşaoğlu. Alternating maximization: Unifying framework for 8 sparse PCA formulations and efficient parallel codes. *arXiv preprint arXiv:1212.4137*, 2012. 17

[40] J. Shao, Y. Wang, X. Deng, and S. Wang. Sparse linear discriminant analysis by thresholding for high dimensional data. *The Annals of Statistics*, 39(2):1241–1265, 2011. 7

[41] H. S. Tapp, M. Defernez, and E. K. Kemsley. Ftir spectroscopy and multivariate analysis can distinguish the geographic origin of extra virgin olive oils. *Journal of agricultural and food chemistry*, 51(21):6110–6115, 2003. 23

[42] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288, 1996. 4

[43] R. Tibshirani, T. Hastie, B. Narasimhan, and G. Chu. Class prediction by nearest shrunken centroids, with applications to dna microarrays. *Statistical Science*, pages 104–117, 2003. 7

[44] M. Welling. Fisher linear discriminant analysis. *Department of Computer Science, University of Toronto*, 2005. 2

[45] D. Witten. *penalizedLDA: Penalized classification using Fisher's linear discriminant*, 2011. R package version 1.0. 18

[46] D. M. Witten and R. Tibshirani. Penalized classification using fisher's linear discriminant. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(5):753–772, 2011. 6, 7, 18, 26

[47] D. M. Witten, R. Tibshirani, and T. Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10(3):515–534, 2009. 17

[48] P. Xu, G. N. Brock, and R. S. Parrish. Modified linear discriminant analysis approaches for classification of high-dimensional microarray data. *Computational Statistics & Data Analysis*, 53(5):1674–1687, 2009. 4

[49] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006. 18

[50] X.-T. Yuan and T. Zhang. Truncated power method for sparse eigenvalue problems. *The Journal of Machine Learning Research*, 14(1):899–925, 2013. 17

[51] Y. Zhang, A. d'Aspremont, and L. El Ghaoui. Sparse PCA: Convex relaxations, algorithms and applications. In *Handbook on Semidefinite, Conic and Polynomial Optimization*, pages 915–940. Springer, 2012. 17

[52] H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005. 6

[53] H. Zou, T. Hastie, and R. Tibshirani. Sparse principal component analysis. *Journal of computational and graphical statistics*, 15(2):265–286, 2006. 17

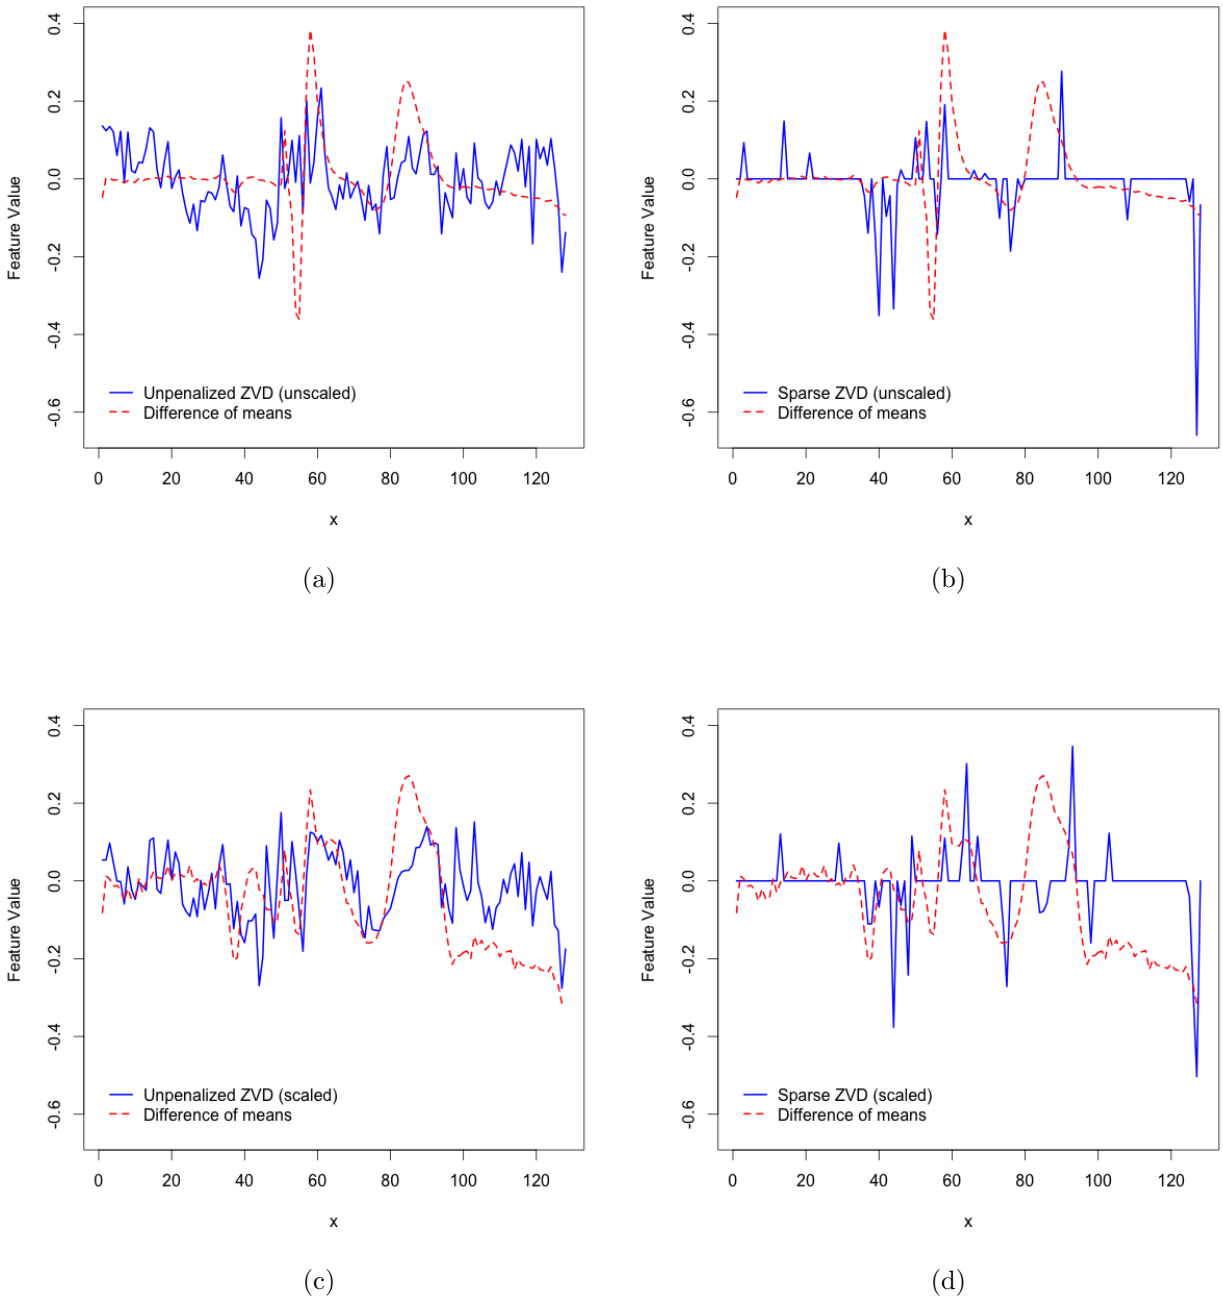

Figure 5.1: Plots of the ZVD (a) and SZVD (b) discriminant vectors with the difference of sample class-mean vectors (the red dashed-line) for the raw *ECGFiveDays* data set; (c) and (d) plot the ZVD and SZVD discriminant vectors after normalizing and centering the *ECGFiveDays* data set. The difference of sample class-mean vectors were rescaled in each plot to emphasize nonzero values of largest magnitude aligning with large differences between the two classes.